Robustness

Robustness

Lars Peter Hansen University of Chicago

Thomas J. Sargent New York University and Hoover Institution

©Lars Peter Hansen and Thomas J. Sargent

12 September 2005

Contents

Acknowledgements

Preface

Part I: Motivation and main ideas

1. Introduction

1.1. Generations of control theory. 1.2. Control theory and rational expectations. 1.3. Misspecification and rational expectations. 1.4. Our extensions of robust control theory. 1.4.1. Discounting. 1.4.2. Representation of worst case shock. 1.4.3. Multiple agent settings. 1.5. Robust control theory, shock serial correlations, and rational expectations. 1.6. Entropy in specification analysis in econometrics. 1.7. Acknowledging misspecification in decision making. 1.8. Organization.

2. Basic ideas and methods

2.1. Introduction. 2.2. Approximating models. 2.2.1. Dynamic programming without model misspecification. 2.3. Measuring model misspecification: entropy. 2.4. Two robust control problems. 2.4.1. Modified certainty equivalence principle. 2.5. Robust linear regulator. 2.6. More general misspecifications. 2.7. A simple algorithm. 2.7.1. Interpretation of the simple algorithm. 2.8. Example: robustness and discounting in a permanent income model. 2.8.1. The LQ permanent income model. 2.8.2. Solution when $\sigma = 0.2.8.3$. Linear regulator for permanent income model. 2.8.4. Effects on consumption of concern about misspecification. 2.8.5. Equivalence of quantities but not continuation values. 2.8.6. Distorted endowment process. 2.8.7. Representing misspecification: a Stackelberg formulation. 2.9. Stabilizing property of shadow price Py_t . 2.10. Forward looking models. 2.11. Concluding remarks. A. Link to a stochastic formulation. B. Matlab programs. vi

vii

Part II: Standard control and filtering

3. Linear control theory

3.1. Introduction. 3.2. Control problems. 3.2.1. Deterministic regulator problem. 3.2.2. Augmented regulator problem. 3.2.3. Discounted stochastic regulator problem. 3.3. Solving the deterministic linear regulator problem. 3.3.1. Nonsingular A_{yy} . 3.3.2. Singular A_{yy} . 3.3.3. Continuous-time systems. 3.4. Computational techniques for solving Riccati equations. 3.4.1. Schur algorithm. 3.4.2. Digression: solving DGE models with distortions. 3.4.3. Doubling algorithm. 3.4.4. Matrix sign algorithm. 3.5. Solving the augmented regulator problem. 3.6. Computational techniques for solving Sylvester equations. 3.6.1. The Hessenberg-Schur algorithm. 3.6.2. Doubling algorithm. 3.7. Concluding remarks.

4. The Kalman filter

4.1. Introduction. 4.2. Review of Kalman filter and preview of main result. 4.2.1. Muth's problem as an example. 4.2.2. The dual to Muth's filtering problem. 4.2.3. The filtering problem. 4.3. Sequence version of primal and dual problems. 4.3.1. Sequence version of Kalman filtering problem. 4.3.2. Sequence version of dual problem. 4.4. Digression: reversing the direction of time. 4.5. Recursive version of dual problem. 4.6. Recursive version of Kalman filtering problem . 4.7. Concluding remarks.

Part III: Robust control and applications

5. Static multiplier and constraint games

5.1. Introduction. 5.2. Phillips curve example. 5.2.1. The government's problem. 5.2.2. Robustness of robust decisions. 5.3. Basic setup with a correct model. 5.4. The constraint game with b = 0.5.5. Multiplier game with b = 0.5.6. The model with $b \neq 0.5.6.1$. Analysis of pathology. 5.7. Probabilistic formulation (b = 0). 5.7.1. Gaussian perturbations. 5.7.2. Letting the minimizing agent make random perturbations when b = 0.5.8. Concluding remarks. A. Rational expectations equilibrium .

59

99

6. Time domain games for robustness

6.1. Several games have identical outcomes. 6.2. The setting. 6.2.1. Constraint and multiplier problems. 6.2.2. Operators and decision rules. 6.3. Timing protocols in three versions of a multiplier game. 6.4. Score card: timing and uses of three multiplier games. 6.5. Game 1: The multiplier game in sequences. 6.5.1. Recursive formulation of maximizing player's problem. 6.5.2. 'Bayesian' interpretation of robust control. 6.5.3. Recursive version of minimizing player's problem. 6.6. Game 2: the Stackelberg multiplier game. 6.7. Game 3: Markov perfect multiplier game. 6.8. Relation between multiplier and constraint problems . 6.9. Recursivity of the constraint game. 6.10. Summary and concluding remarks. A. Certainty equivalence. B. Useful formulas. 6.B.1. A single Riccati equation. 6.B.2. Robustness bound. 6.B.3. Special case: a pure forecasting problem. C. Details of a proof of Theorem 6.5.2.

7. Frequency domain games and criteria for robustness

7.1. Three frequency domain criteria indexed by (x_0, η) . 7.2. The Stackelberg game in the time domain. 7.3. Stackelberg game in frequency domain. 7.3.1. Fourier transforms. 7.3.2. H_2 criterion. 7.3.3. The Stackelberg game in the frequency game. 7.3.4. Version 2: the H_{∞} Criterion. 7.4. A multiplier game for version 3. 7.5. The multiplier problem and the entropy criterion. 7.5.1. A robustness bound. 7.5.2. Entropy is the indirect utility function of the multiplier problem. 7.6. Relation of Stackelberg multiplier game to entropy criterion. 7.7. Etymology of 'entropy'. 7.8. Risk sensitivity. 7.9. Risk aversion across frequencies. 7.10. Concluding remarks. A. Infimization of H_{∞} . B. A dual prediction problem. C. Duality. 7.C.1. Evaluating a given control law. 7.C.2. When $\theta = \theta_F$. 7.C.3. Failure of entropy condition. D. Proof of theorem 7.5.4. E. Stochastic interpretation of H_2 . 7.E.1. Stochastic counterpart.

8. Calibrating θ with detection probabilities

8.1. The role of randomness. 8.1.1. Approximating and distorting models. 8.2. Detection error probabilities. 8.3. Details. 8.3.1. Likelihood ratio under the approximating model. 8.3.2. Likelihood ratio under the distorted model. 8.3.3. The detection error probability. 8.3.4. Ball's model. 8.3.5. Robustness in a simple macroeconomic model.

138

177

9. A permanent income model

9.1. Introduction. 9.2. A robust permanent income theory. 9.3. Solution when $\sigma = 0.9.3.1$. The $\sigma = 0$ benchmark case. 9.3.2. Observational equivalence (for quantities) of $\sigma = 0$ and $\sigma \neq 0.9.3.3$. Observational equivalence: intuition. 9.3.4. Observational equivalence: formal argument. 9.3.5. Precautionary savings interpretation. 9.4. Observational equivalence and distorted expectations. 9.4.1. Distorted endowment process. 9.5. Another view of precautionary savings. 9.6. Frequency domain representation. 9.7. Detection error probabilities. 9.8. Robustness of decision rules. 9.9. Concluding remarks. A. Parameter values. B. Another observational equivalence result.

10. Competitive equilibrium models

10.1. Introduction. 10.2. Pricing risky claims. 10.3. Types of competitive equilibria. 10.4. Information, preferences, and technology. 10.4.1. Information. 10.4.2. Preferences. 10.4.3. Technology. 10.4.4. Planning problem. 10.4.5. Imposing stability. 10.5. Arrow-Debreu equilibrium. 10.5.1. The price system at time 0. 10.5.2. The household. 10.5.3. The firm. 10.5.4. Competitive equilibrium with time-zero trading. 10.5.5. Equilibrium computation. 10.5.6. Shadow prices. 10.5.7. Recursive representation of time 0 prices. 10.5.8. Recursive representation of household's problem. 10.5.9. Units of prices and reopening markets. 10.6. Sequential markets with Arrow securities. 10.6.1. Arrow securities. 10.6.2. Recursive competitive equilibrium. 10.7. Asset pricing in a nutshell. 10.8. Partial equilibrium interpretation. 10.9. Concluding remarks.

11. Competitive equilibrium under robustness

11.1. Introduction. 11.2. A pure endowment economy. 11.2.1. The planning problem. 11.2.2. Household problem. 11.3. A robust planning problem. 11.4. Min-max representation of the household's problem in an Arrow-Debreu equilibrium. 11.4.1. Sequence problem of maximizing player. 11.4.2. Digression about computing μ_0^w . 11.4.3. Sequence problem of minimizing player. 11.5. A decentralization with Arrow securities. 11.5.1. A robust consumer trading Arrow securities. 11.5.2. The inner problem. 11.6.3. The outer problem. 11.6. An expost Bayesian planning problem. 11.6.1. Remarks on practicality. 11.7. Two asset pricing strategies. 11.7.1. Pricing from the robust planning problem. 11.8. Concluding remarks. A. Decentralization of partial equilibrium.

273

Contents

12. Asset pricing

12.1. Introduction. 12.2. Approximating and distorted models. 12.3. Asset pricing without robustness. 12.4. Asset pricing with robustness. 12.4.1. Adjustment of stochastic discount factor for fear of model misspecification. 12.4.2. Reopening markets. 12.5. Pricing single period payoffs. 12.5.1. Calibrated market prices of Knightian uncertainty. 12.6. A model of occupational choice and pay. 12.6.1. A one-occupation model. 12.6.2. Equilibrium with no concern about robustness. 12.6.3. Example. 12.7. Concluding remarks. A. Solving Ryoo and Rosen's model by hand.

Part IV: Robust filtering

13. A robust filtering problem

13.1. Filtering. 13.1.1. Warning about recycling of notation. 13.2. Summary of robust filtering and duality. 13.2.1. Duality of ordinary filtering and control. 13.2.2. A robust linear regulator. 13.2.3. A dual robust filtering problem. 13.2.4. The robust filter when Hx_{-t} is observable. 13.3. The robust filtering problem. 13.3.1. The evil agent's problem. 13.3.2. The dual to the evil agent's problem. 13.3.3. Computing K. 13.3.4. Matlab programs. 13.4. Robustifying Muth's filter. 13.4.1. Ordinary Kalman filter. 13.4.2. Illustrations. 13.5. Another example. 13.6. Discounting and the direction of time. 13.7. Another perturbation. 13.7.1. Transition to next chapter. A. Another way to compute the worst-case shock. B. A Bayesian interpretation of the robust K.

14. Joint control and estimation

14.1. Introduction. 14.2. A recursive control and filtering problem. 14.2.1. The decision maker's approximating model. 14.2.2. Two sources of statistical perturbation. 14.2.3. Two operators. 14.2.4. A recursive formulation for control and estimation. 14.2.5. A certainty equivalent short cut. 14.2.6. Computing the \mathbf{T}^1 operator. 14.2.7. Worst case signal distribution. 14.3. Examples. 14.3.1. Asset pricing with robustness and filtering. A. Worst case signal distribution. 339

315

Part V: More applications

15. Multiple agents

15.1. Introduction. 15.2. Markov perfect equilibria with robustness. 15.2.1. Computational algorithm: iterating on stacked Bellman equations. 15.2.2. Bayesian interpretation and belief heterogeneity. 15.3. Concluding remarks.

16. Robustness in forward looking models

16.1. Introduction. 16.1.1. Related literature. 16.2. The robust Stackelberg problem. 16.2.1. Multiplier version of the robust Stackelberg problem. 16.3. Solving the robust Stackelberg problem. 16.3.1. Step 1: solve a robust linear regulator. 16.3.2. Step 2: use the stabilizing properties of shadow price Py_t . 16.3.3. Key insight. 16.3.4. Step 3: convert implementation multipliers into state variables. 16.4. Incorporating robustness for the followers. 16.4.1. An approach enabled by the Bellman-Isaacs condition. 16.4.2. Heterogeneity of worst case models. 16.4.3. Evolution under approximating model and robust rules. 16.4.4. Evolution of the economy under Stackelberg solution. 16.5. A monopolist with a competitive fringe. 16.5.1. The competitive fringe. 16.5.2. The monopolist's problem. 16.5.3. Recursive representation of a competitive firm's problem. 16.5.4. Representation of the monopolist's decision rule under the approximating model. 16.5.5. Numerical example. 16.6. Concluding remarks. A. Invariant subpace method. B. The Riccati equation. 16.B.1. The Riccati equation. C. Another Bellman equation.

17. Non-linear models

17.1. Introduction. 17.1.1. The \mathcal{R} operator for LQ problems. 17.1.2. Markov perturbations. 17.1.3. Relative entropy. 17.2. Value function for robustness. 17.2.1. Gaussian example. 17.3. Large deviation interpretation of \mathcal{R} .

359

Contents

18. Ramsey plans

18.1. Introduction. 18.1.1. Exogenous processes and information. 18.1.2. Technology. 18.1.3. Representation of price system. 18.1.4. Households. 18.1.5. Government. 18.1.6. Equilibrium. 18.1.7. Properties. 18.1.8. Ramsey problem. 18.1.9. Solution strategy. 18.1.10. Computation with no concern about robustness. 18.1.11. Key idea. 18.1.12. Execution. 18.2. Modifications for robustness. 18.2.1. Main idea. 18.2.2. Assumption 1'. 18.2.3. Assumption 2'. 18.2.4. Ramsey planner's worst-case model. 18.3. Computations. 18.4. Simulations at approximating model. 18.5. Computed Ramsey plans and allocations. 18.5.1. Markov case. 18.5.2. Stochastic difference equation.

19.	References 4	27
20.	Index 4	39

- 21. Author Index443
- 22. Matlab Index 445

Acknowledgments

For criticisms of previous drafts and stimulating discussions of many issues we thank Fernando Alvarez, Luca Benati, Dirk Bergemann, V.V. Chari, Jack Y. Favilukis, Patrick Kehoe, Francesco Lippi, Pascal Maenhout, Ricardo Mayer, Tomasz Piskorski, Joseph Pearlman, Mark Salmon, Martin Schneider, Joseph Teicher, Aaron Tornell, Peter von zur Muehlen, Neng Wang, Pierre-Olivier Weill, and François Velde. In addition to providing comments, Yong Shin, Stijn Van Nieuwerburgh, Chao D. Wei, and Mark Wright helped with the computations. Our editor at Princeton University Press, Peter Dougherty, provided encouragement and valuable suggestions about style and presentation. We thank John Doyle, an accomplished artist as well as a creator of robust control theory, for letting us reproduce Fig. 1.1.1, whose disturbing image of a robust control theorist for us conjures up why Arthur Goldberger and Robert E. Lucas, Jr., warned us to beware of a theorist bearing a free parameter, in this case θ .

Preface

A profitable decision rule for us has been, 'if Peter Whittle wrote it, read it.' Whittle's 1963 book Prediction and Regulation by Linear Least Squares Methods (reprinted and revised in 1983), taught early users of rational expectations econometrics, including ourselves, the classical time series techniques that were perfect for putting the idea of rational expectations to work. When we became aware of Whittle's 1990 book, Risk Sensitive Control and later his 1996 book Optimal Control: Basics and Beyond, we eagerly worked our ways through them. These and other books on robust control theory like Başar and Bernhard's H^{∞} -Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach provide tools for rigorously treating the 'sloppy' subject of how to make decisions when one does not fully trust a model and open the possibility of rigorously analyzing how wise agents should cope with fear of misspecification. While Whittle mentioned a few economic examples, the methods that he and other authors of robust and risk sensitive control theories had developed were mainly designed for types of problems that differ in some significant ways from economic problems. Therefore, we soon recognized that risk sensitive and robust control methods would require some modifications if they were to be applied successfully to economic problems. That is why we started the research that led to this book. We do not claim to have attained a general theory of how to make economic decisions while acknowledging misspecification, but only to have begun to study this difficult and important problem that has concerned every researcher who has estimated a rational expectations model, every central banker who has used dubious models to guide his monetary policy decisions, and every macroeconomist who has ever viewed formal estimation as wrongheaded and instead 'calibrated' the parameters of a complete, but admittedly highly stylized, model.

– xiii –

Part I Motivation and main ideas

Chapter 1 Introduction

1.1. Generations of control theory

Figure 1.1.1 reproduces John Doyle's cartoon of developments in optimal control theory since World War II.¹ Two scientists in the upper panels use different mathematical methods to devise control laws and estimators. The person on the left uses classical methods (Euler equations, z-transforms, lag operators) and the one on the right uses modern recursive methods (Bellman equations, Kalman filtering). The scientists in both top panels completely trust their models of the transition dynamics and use mathematical procedures that assume that those models are true. The gentleman in the lower panel shares the objectives of his predecessors from the 50s, 60s, and 70s, but regards his model only as an approximation to an unknown and unspecified model that he thinks actually generates the data. He seeks decision rules and estimators that work over a continuous set of models near his approximating model. The H_{∞} in his postmodern tatoo and the θ on his staff are alternative ways to express his doubts about his model by measuring the discrepancy of the true data generating mechanism from his approximating model. As we shall learn in later chapters, the parameter θ is interpretable as a penalty on a measure of discrepancy (the entropy) between his approximating model and the model that actually generates the data. The H_{∞} refers to a measure of his objective function under a greatest lower bound on the multiplier θ .

- 1 -

¹ John Doyle consented to let us reproduce this drawing, which appears in Zhou, Doyle, and Glover (1996). We have changed Doyle's notation by making θ (Doyle's μ) the free parameter borne by the post-modern control theorist.



Figure 1.1.1: A pictorial history of control theory (courtesy John Doyle). Beware of theorists bearing a free parameter, θ .

1.2. Control theory and rational expectations

Classical and modern control theory supplied perfect tools for applying Muth's (1961) concept of rational expectations to a variety of problems in dynamic economics. Rational expectations initially diffused slowly after Muth's (1961) paper precisely because in 1961 most economists were not sufficiently familiar with the tools lampooned in the top panel of Fig. 1.1.1. Rational expectations took hold in the 1970's only after a generation of young macroeconomists had learned those tools. Ever since, macroeconomists and rational expectations

econometricians have gathered inspiration and ideas from classical and recursive control theory. 2

Ironically, when macroeconomists began to apply classical and modern control theory in the late 1970's, control theorists and applied mathematicians were trying to find ways to relax the assumption that the decision maker trusts his model. They sought new control and estimation methods to repair adverse outcomes they had experienced from applying classical and modern control theory to a variety of engineering and physical problems. They thought that model misspecification explained why actual outcomes were sometimes much worse than control theory predicted and therefore sought controls and estimators that acknowledged model misspecification. That is how robust control and estimation theory came into being.

1.3. Misspecification and rational expectations

It is an understatement to say that model misspecification is at least as much of a problem in economics as it is in physics and engineering. This book borrows, adapts, and extends tools from the literature on robust control and estimation to model decision makers who regard their models as approximations. A decision maker believes that the data come from an unknown member of a *set* of unspecified models near his approximating model.³ Concern about model misspecification induces a decision maker to want decision rules that work over that set of nearby models.

If they lived inside rational expectations models, decision makers would not have to worry about model misspecification. They could trust their model because subjective and objective probability distributions (i.e., models) coincide. Rational expectations theorizing removes agents' personal models as elements of the model.⁴

 $^{^2\,}$ See Stokey and Lucas with Prescott (1989), Ljungqvist and Sargent (2004), and Hansen and Sargent (1991) for many examples.

 $^{^{3}}$ We say unspecified because they are formulated as vague perturbations to the decision maker's approximating model.

⁴ In a rational expectations model, each agent's model (i.e., his subjective joint probability distribution over exogenous and endogenous variables) is an equilibrium outcome, not something to be specified by the model builder. Its early advocates in econometrics emphasized the empirical power that followed from the fact that the rational expectations hypothesis

Introduction

Although the artificial agents within a rational expectations model trust the model, a model's author often doubts it, especially after performing specification tests or when calibrating it. There are several good reasons for wanting to extend rational expectations models to acknowledge fear of model misspecification.⁵ First, doing so accepts Muth's (1961) intention of putting econometricians and the agents being modelled on the same footing: because econometricians face specification doubts, the agents inside the model might too.⁶ Second, in various contexts, rational expectations models under-predict prices for risk that are revealed by asset market data. For example, relative to standard rational expectations models, actual asset markets seem to assign prices to macroeconomic risks that are too high. The equity premium puzzle is one manifestation of this is mispricing.⁷ Agents' caution in responding to concerns about model misspecification can raise the theoretical values of the prices to be assigned to those macroeconomic risks. This reason for studying robust decisions is positive and is to be judged by how it helps explain market data. A third reason for studying the robustness of decision rules to model misspecification is normative. A long tradition dating back to Friedman (1953), Bailey (1971), Brainard (1967) and Sims (1971, 1972) advocates framing macroeconomic policy rules and interpreting econometric findings in light of doubts about model specification, though how those doubts have been formalized in practice has varied.

eliminates all free parameters associated with peoples' beliefs. For example, see Hansen and Sargent (1980) and Sargent (1981).

 $^{^{5}}$ In chapter 16, we explore various mappings, the fixed points of which can be used to restrict a robust decision makers' approximating model. As is usually the case with rational expectations models, we are silent about the process by which an agent arrives at his approximating model. A qualification to the claim that rational expectations models do not model the process by which agents' model is formed comes from the literature on learning, in which agents who use recursive least squares learning schemes eventually come to have rational expectations. Early examples of such work are Bray (1982), Marcet and Sargent (1989), and Woodford (1990). See Evans and Honkapohja (2001) for new results.

⁶ This argument might offend someone with a preference against justifying modelling assumptions on behavioral grounds.

⁷ A related finding is that rational expectations models impute low costs to business cycles. See Hansen, Sargent, and Tallarini (1999), Tallarini (2000), and Alvarez and Jermann (1999).

1.4. Our extensions of robust control theory

When we say that we 'adapt and extend' robust control theory for economic problems, we have three main things in mind: discounting, a reinterpretation of the 'worst case shock process', and equilibrium reasoning.

1.4.1. Discounting

Most treatments of robustness in the control theory treat undiscounted problems, and the few formulations of discounting that do appear differ from the way macroeconomics would set things up.⁸ In this book, we formulate discounted problems that preserve the recursive structure of decision problems that macroeconomists and other applied economists use so widely.

1.4.2. Representation of worst case shock

As we shall see, in existing formulations of robust control theory, the misspecification shocks and the worst case shock process, in particular, is allowed to feed back on endogenous state variables that are influenced by the decision maker, an outcome that in some contexts appears to confront the decision maker with peculiar incentives to manipulate future values of some of those shocks by adjusting his current decisions. Some critics of robust control theory in economics like Christopher Sims have questioned the plausibility of the notion that the decision maker is concerned about *any* misspecifications that can be represented in terms of shocks that feed back on state variables under his partial control. In chapter 6, we use a device informally known as the 'Big K, little k trick' from the literature on recursive competitive equilibria to reformulate perturbations to an approximating model, and the worst case shock in particular, as exogenous processes that cannot be influenced by the decision maker. As we illustrate in the analysis of the permanent income model of chapter 9, this reinterpretation of the worst case shock process is useful in a variety of economic models.

⁸ See Whittle (1990) and Hansen and Sargent (2005XXX-commitment).

1.4.3. Multiple agent settings

In formulations from the control theory literature, a description of the decision maker's model of the state transition law is a primitive element of the statement of the problem. In multi-agent dynamic economic problems, it is not. Instead, the parts of the decision maker's transition law governing endogenous state variables, such as aggregate capital stocks, are affected by other agents' choices and therefore are outcomes of equilibrium reasoning. In this book, we describe ways of formulating the decision maker's approximating model when he and other decision makers are all concerned about model misspecification, perhaps in differing degrees. We proceed in a methodologically conservative way by imposing a common approximating model on all decision makers, but allow them to express different degrees of mistrust of that model. As we explain in chapters 15, 11, 16, and 18, this is a conservative procedure in the sense that it adapts the concept of a Nash equilibrium to include concerns about robustness.

1.5. Robust control theory, shock serial correlations, and rational expectations

Ordinary optimal control theory assumes that decision makers know the model in the form of a transition law linking the motion of state variables to controls. The optimization problem associates a distinct decision rule with each specification of shock processes. Many aspects of rational expectations models flow from this association.⁹ For example, the cornerstone of the Lucas (1976) Critique is the finding that, under rational expectations, decision *rules* are functionals of the serial correlations of shocks. Rational expectations econometrics achieves parameter identification by exploiting the structure of the function that maps shock serial correlation properties to decision rules.¹⁰

Robust control theory alters the mapping from shock temporal properties to decision rules. Robust control theory treats the decision maker's model as an approximation and seeks a *single* rule to use for a *set* of models that might govern the data. The alternative models are specified vaguely in terms of distortions to

 $^{^{9}\,}$ Stoke and Lucas with Prescott (1989) is a standard reference on using control theory to construct dynamic models in macroeconomics.

¹⁰ See Hansen and Sargent (1980, 1981, 1991).

the shock processes in the decision maker's model. Because they are allowed to feed back arbitrarily on the history of the states, such distortions can represent general misspecifications of the dynamics.

As emphasized by Hansen and Sargent (1980, 1981, 1991), the econometric content of the rational expectations hypothesis is a set of cross-equation restrictions that cause particular combinations of parameters characterizing the stochastic processes impinging on agents' constraint sets to appear in decision rules. A concern for model misspecification alters these cross-equation restrictions by inspiring the robust decision maker to *twist* or *slant* those stochastic processes in ways designed to make his decision rule less fragile to misspecification. Formulas presented in chapters 2 and 6 imply that the Hansen-Sargent (1980, 1981) formulas for those cross-equation restrictions also describe the behavior of the robust decision maker, provided that we use the appropriated *slanted* laws of motion in the Hansen-Sargent forecasting formulas. This finding is interesting because it shows how robust control theory adds a concern about misspecification in a way that preserves the econometric discipline imposed by rational expectations econometrics.

1.6. Entropy in specification analysis in econometrics

The statistical and econometric literatures on the analysis of model misspecification supply tools for measuring discrepancies between models and for thinking about decision making in the presence of model misspecification.

Where y^* denotes next period's value of a state vector y, let the data truly come from a Markov process with one step transition density $f(y^*|y)$ for a Markov process that we assume has invariant distribution $\mu(y)$. is the invariant Let the econometrician's model be $f_{\alpha}(y^*|y)$ where $\alpha \in A$ and A is a compact set of values for a parameter vector α . If there is no $\alpha \in A$ such that $f_{\alpha} = f$, we say that the econometrician's model is *misspecified*. Assume that the econometrician estimates α by maximum likelihood. Under some regularity conditions, the maximum likelihood estimator $\hat{\alpha}_{\rho}$ converges in large samples $\mathrm{to}^{\,11}$

$$\operatorname{plim} \hat{\alpha}_{o} = \operatorname{argmin}_{\alpha \in A} \int I(\alpha, f)(y) \, d\,\mu(y) \equiv I(f) \tag{1.6.1}$$

where $I(\alpha, f)(y)$ is the conditional relative entropy of model f with respect to model f_{α} defined as the expected value of the logarithm of the likelihood ratio evaluated with respect to the true conditional density $f(y^*|y)$

$$I(\alpha, f)(y) = \int \log\left(\frac{f(y^*|y)}{f_{\alpha}(y^*|y)}\right) f(y^*|y) \, dy^*.$$
(1.6.2)

The quantity $I(\alpha, f)(y)$ is called the conditional entropy of model f_{α} relative to model f. It can be shown that $I(\alpha, f)(y) \geq 0$. When the model is misspecified, the minimized value of relative entropy defined on the right side of (1.6.1) is positive. Figure 1.6.1 depicts how the probability limit of the estimator of the parameters of a misspecified model α_o makes I(f)(y) as small as possible.



Figure 1.6.1: Econometric specification analysis. Suppose that the data generating mechanism is f and that the econometrician fits a parametric class of models $f_{\alpha} \in A$ to the data and that $f \notin A$. Maximum likelihood estimates of α eventually select the misspecified model the model f_{α_o} that is closest to f as measured by entropy I(f).

Sims (1993) and Hansen and Sargent (1993) have used the preceding analysis of estimation of misspecified models has been used to deduce the consequences of various types of misspecification for estimates of particular parameters of dynamic stochastic models.¹² For example, they studied the consequences of using

¹¹ Versions of this result occur in White (1982), Vuong (1989), Sims (1993), Hansen and Sargent (1993), and Gelman et. al. (XXXX).

¹² Also see Vuong (1989).

seasonally adjusted data to estimate rational expectations models of decisions makers who actually base their decisions on raw seasonally unadjusted data.

1.7. Acknowledging misspecification in decision making

To study decision making in the presence of model misspecification, we turn this analysis on its head by taking $f_{\alpha_{\alpha}}$ as a given unreliable approximating model and surrounding it with a set of unknown possible data generating processes, one unknown element of which is the true process f. See figure 1.7.1. Decision makers know their approximating model $f_{\alpha_{\alpha}}$, but because they don't know f, they must base their decisions on the only explicitly specified model available, namely, the misspecified f_{α_o} . We say nothing about the process of discovery that culminated in the decision maker's approximating model $f_{\alpha_o}(y^*|y)$.¹³ We also take for granted the decision maker's parameter estimates α_o .¹⁴ We impute to the decision maker some doubts about his model. In particular, the decision maker suspects that the data are actually generated by another model $f(y^*|y)$ with relative entropy $I(\alpha_o, f)(y)$. The decision maker thinks that his model is a good approximation in the sense that $I(\alpha_o, f)(y)$ is not too large, and wants to make decisions that will be good when $f \neq f_{\alpha_o}$. We endow the decision maker with a discount factor β and construct the following intertemporal measure of model misspecification:¹⁵

$$I = E_f \sum_{t=0}^{\infty} \beta^t I(\alpha_o, f)(y_t)$$

where E_f is the mathematical expectation evaluated with respect to the distribution f. Our decision maker confronts model misspecification by seeking a decision rule that will work well across a set of models for which $I \leq \tilde{\eta}_0$, where $\tilde{\eta}_0$ measures the set of models F surrounding his approximating model

 $^{^{13}}$ Tom: here is where we should cite Kreps's chapter 11.

¹⁴ In chapter 8, we entertain the hypothesis that the decision maker has estimated his model by maximum likelihood using a data set of length T and use Bayesian detection error probabilities to guide the choice of a set of models against which the decision maker wants to be robust.

¹⁵ Hansen and Sargent (2005a, 2005b XXXX) provide an extensive discussion of reasons for adopting this measure of model misspecification.

 f_{α} . Fig. 1.7.1 portrays the decision maker's view of the world. The decision maker wants a single decision rule that will work reliably for all models in the set displayed in Fig. 1.7.1.¹⁶ This book describes how he can form such a robust decision rule by using a Bellman equation in which he maximizes his intertemporal objective over decision rules while a hypothetical malevolent nature minimizes that same objective by choosing a model f. Of course, this hypothetical malevolent nature is just a device that the decision maker uses to construct a lower bound on the performance of his rule. A decision maker interested in robustness is naturally interested in constructing bounds on the performance of a decision rule, and the malevolent agent assists him in that endeavor.



Figure 1.7.1: Robust decision making: A decision maker with model f_{α_o} suspects that the data are actually generated by a nearby model f, where $I(\alpha_o, f) \leq \eta$.

¹⁶ Note how this uproots the one-to-one mapping from transition laws f to decision rules that is emphasized in the Lucas (1976) critique.

1.8. Organization

This monograph displays alternative ways to express and respond to a decision maker's doubts about model specification. We study both control and estimation (or filtering) problems, and both single- and multiple-agent settings. As already mentioned, we adapt and extend results from the robust control literature in two important ways. First, while the control literature focuses on undiscounted problems, we formulate discounted problems. Incorporating discounting involves substantial work, especially in chapter 7, and requires paying special attention to initial conditions. Second, we analyze three types of economic environments with multiple decision makers who are concerned about model misspecification: (1) a competitive equilibrium with complete markets in history-date contingent claims and a representative agent who fears model misspecification (chapters 11 and 12); (2) a Markov perfect equilibrium of a dynamic game with multiple decision makers who fear model misspecification (chapter 15); and (3) a Stackelberg or Ramsey problem in which both the leader and the followers fear model misspecification (chapters 16 and 18). Thinking about model misspecification in these environments requires that we introduce an equilibrium concept that extends rational expectations. We stay mostly, but not exclusively, within a linear quadratic framework (see chapter 17 for the more general case), in which a pervasive certainty equivalence principle allows a nonstochastic presentation of most of the control and filtering theory.

The monograph is organized as follows. Chapter 2 summarizes a set of practical results at the lowest possible technical level. A message of this chapter is that, although sophisticated arguments from chapters 6 and 7 are needed fully to justify the techniques of robust control, the techniques themselves are as easy to apply as the ordinary dynamic programming techniques that are now widely used throughout macroeconomics and applied general equilibrium theory. Chapter 2 uses linear quadratic dynamic problems to convey this message, but the message applies more generally, as we shall illustrate in chapters 17 and 18.

Chapters 3 and 4 are about optimal control and filtering when the decision maker trusts his model. Chapter 3 sets forth important principles by summarizing results about the classic optimal linear regulator problem. This chapter builds on the survey by Anderson, Hansen, McGrattan, and Sargent (1996) and culminates in a description of invariant subspace methods for solving linear optimal control and filtering problems and also for solving dynamic linear

Introduction

equilibrium models. Later chapters apply these methods to various problems: to compute robust decision rules as solutions of zero-sum two-player games; to compute robust filters via another zero-sum two-player game; and to compute equilibria of robust Stackelberg or Ramsey problems in macroeconomics.

Chapter 4 emphasizes that the Kalman filter is the *dual* (in a sense familiar to economists from their use of Lagrange multipliers) of the basic linearquadratic dynamic programming problem of chapter 3. We exploit duality relations often in subsequent chapters.

Within a one-period setting, chapter 5 introduces two-person zero-sum games as a way to induce robust decisions. Although the forms of model misspecifications considered in this chapter are very simple relative to those considered in subsequent chapters, the static setting of chapter 5 is a good one for addressing some important conceptual issues. In particular, in this chapter for the first time we state multiplier and constraint problems, alternative optimization problems that induce robust decision rules. We use the Lagrange multiplier theorem to show the connection between the two problems.

Chapters 6 and 7 extend and modify results in the control literature to formulate robust control problems with discounted quadratic objective functions and linear transition laws. Incorporating discounting requires carefully restating the control problems used to induce robust decision rules. Chapters 6 and 7 describe two ways to alter the discounted linear quadratic optimal control problem in a way to induce robust decision rules: (1) to form one of several two-player, zero-sum games in which nature chooses from a set of models in a way that makes the decision maker want robust decision rules; and (2) to adjust the continuation value function in the dynamic program in a way that encodes the decision maker's preference for a robust rule. The continuation value that works comes from the minimization piece of one of the zero-sum two-player games in (1). In category (1), we present a detailed account of several two-person zero-sum games with different timing protocols, each of which induces a robust decision rule. As an extension of category (2), we present three specifications of preferences that embed a preference for robust rules. Two of them are expressed in the frequency domain: the H_{∞} and entropy criteria. The entropy objective function summarizes model specification doubts with a single parameter. We describe how that parameter relates to a Lagrange multiplier in a two-player zero-sum game, and also to the risk-sensitivity parameter of Jacobson (1973)

Organization

and Whittle (1990), as modified for discounting by Hansen and Sargent (1995).

Chapters 6 and 7 show how robustness is induced by using min-max strategies: the decision maker maximizes while nature minimizes over a set of models that are close to the approximating model. Chapter 8 gives our method for calibrating a measure of proximity based on statistical detection theory.¹⁷ There are alternative timing protocols in terms of which a zero-sum two player game can be cast. A main finding of chapter 6 is that zero-sum games that make a variety of different timing protocols share outcomes and representations of equilibrium strategies. This important result lets us use recursive methods to compute our robust rules and it facilitates important analytical approaches for computing equilibria in multiple agent economics.

Arthur Goldberger and Robert E. Lucas, Jr., warned applied economists to beware of theorists bearing free parameters (see Fig. 1.1.1). Relative to settings in which decision makers completely trust their models, the multiplier and the constraint problems of chapters 6 and 7 each bring one new free parameter that expresses a concern about model misspecification, θ for the multiplier problem and η for the constraint problem. Each of these parameters measures sets of models near the approximating model against which the decision maker seeks a robust rule. Chapter 8 proposes a way to calibrate these parameters by using the statistical theory for discriminating models. We apply this theory in chapter 9.

Chapter 9 uses the permanent income model of consumption as a laboratory for illustrating some of the concepts from chapters 6 and 7. Because he prefers smooth consumption paths, the permanent income consumer's savings are designed to attenuate the effects of income fluctuations on his consumption. A robust consumer engages in a kind of precautionary savings because he suspects error in the specification of the income process. We will also use the model of chapter 9 as a laboratory for asset pricing in chapter 12. But first, chapters 10 and 11 describe how to decentralize the solution of a planning problem with a competitive equilibrium. Chapter 10 sets out a class of dynamic economies and describes two decentralizations, one with trading of history-date contingent commodities once and for all at time zero, and another with sequential trading of one-period Arrow securities. In that sequential setting, we give a recursive representation of equilibrium prices. Chapter 10 describes a setting

¹⁷ See Anderson, Hansen, and Sargent (2003) for a further discussion.

Introduction

where the representative agent has no concern about model misspecification, while chapter 11 extends the chapter 10 characterizations to situations where the representative decision maker fears model misspecification.

Chapter 12 builds on the chapter 11 results to show how fear of model misspecification affects asset pricing. We show how, from the vantage point of the approximating model, a concern for robustness induces a multiplicative adjustment to the stochastic discount factor, where the adjustment measures fear that the approximating model is misspecified. The adjustment for robustness mathematically resembles similar adjustments that finance experts used to construct risk neutral probability measures for pricing assets. We describe the basic theory within a class of linear quadratic general equilibrium models and then a calibrated version of the permanent income model of chapter 9.

Chapter 13 extends the analysis of filtering from chapter 4 by describing a discounted robust filtering problem that is dual to the control problem of chapter 6. We discover this problem by stating and solving a conjugate problem of a kind familiar to economists through duality theory. By faithfully following where duality leads us, we discover a filtering problem that is peculiar (but not necessarily uninteresting) from an economic stand point. There are two peculiarities. First, the decision maker discounts the more distant past. Second, 'bygones are not bygones': the decision makers concerns about *past* returns affect his current estimate of a hidden state vector.

Chapter 14 studies robust filtering again and, by using a different criterion than chapter 13, finds a different robust filter. We argue that the chapter 14 filter is the appropriate one for many problems and give some examples. The different filter that emerge from chapters 13 and 14 illustrate how robust decision rules are 'context specific' in the sense that they depend on the common objective function in the two-player zero-sum game that is used to induce a robust decision rule. This theme will run through our monograph. Following Hansen, Sargent, and Wang (2001), we use the permanent income model of chapter 9 as a laboratory for illustrating a filtering problem.

Chapters 15 and 16 describe two more settings with multiple decision makers and introduce an equilibrium concept that extends rational expectations in what we think is a natural way. In a rational expectations equilibrium, all decision makers completely trust a common model and important aspects of that model, those governing endogenous state variables, are equilibrium outcomes.

Organization

The source of the powerful cross-equation restrictions that are the hallmark of rational expectations econometrics is that decision makers share a common model and that model is presumed to govern the data.¹⁸ To preserve that empirical power in our equilibria with multiple decision makers who fear model misspecification, we impose that all decision makers share a common approximating model.¹⁹ The pieces of that model that describe endogenous state variables are equilibrium outcomes that depend on agents' robust decision making processes, and in particular, on how their min-max problems.

Chapter 15 describes how to implement this equilibrium concept in the context of a two-player dynamic game in which the players share a common approximating model and in which each player makes robust decisions by solving a two-player zero-sum game taking the approximating model as given. We show how to compute the approximating model by appropriately stacking and solving robust versions of the Bellman equations and first-order conditions for the two decision makers. While the equilibrium imposes a common approximating model, the worst-case models of the two decision makers diverge because their objectives diverge. In this restricted sense, the model produces a version of endogenous heterogeneity of beliefs.

In chapter 16, we alter the timing protocol in a way that invites one decision maker to manipulate the worst-case beliefs of the other decision makers and thereby affect their decision rules. Chapters 16 and 18 study versions of a macroeconomic control problem, called a Ramsey problem, where a leader wants optimally to control followers who are forecasting the leader's controls. We describe how to compute a robust Stackelberg policy when the Stackelberg leader can commit to a rule. We accomplish that by using a robust version of the optimal linear regulator or else one of the invariant subspace methods of chapter 3.

Chapter 17 tells how the key ideas about robustness generalize to models that are not linear-quadratic. Then chapter 16 modifies Lucas and Stokey's

¹⁸ Lars: you might want to edit this footnote. The restriction that they share a common model is the feature that makes free parameters governing expectations disappear. The restriction that that common model actually governs the data is what legitimizes a law of large numbers that underlies rational expectations econometrics.

¹⁹ In the empirical applications of Hansen, Sargent, and Tallarini (1999) and Anderson, Hansen, and Sargent (2003), we also maintain the second aspect of rational expectations modelling, namely, that the decision makers' approximating model actually *does* generate the data.

Introduction

model of optimal taxation by allowing the representative consumer to doubt the specification of the Markov process for government expenditures. That alteration gives the Ramsey planner a motive to manipulate the representative consumer's worst case model as part of the process by which the planner manipulates equilibrium prices in solving the Ramsey problem.

Chapter 2 Basic ideas and methods

There are two different drives toward exactitude that will never attain complete fulfillment, one because "natural" languages always say something more than formalized languages can – natural languages always involve a certain amount of noise that impinges on the essentiality of the information – and the other because, in representing the density and continuity of the world around us, language is revealed as defective and fragmentary, always saying something less with respect to the sum of what can be experienced. — Italo Calvino, Six Memos for the Next Millenium, 1996, pp. 74-75

2.1. Introduction

Standard control theory tells a decision maker how to make optimal decisions when his model is correct. Robust control theory tells him how to make good decisions when his model approximates the correct one. This chapter summarizes methods for computing robust decision rules when the decision maker's criterion function is quadratic and his approximating model is linear.¹ We modify the Bellman equation and the Riccati equation associated with the standard linear-quadratic dynamic programming problem to incorporate concerns about misspecification of the transition law. The adjustments to the Bellman equation have alternative representations, each of which has practical uses in various contexts.² This chapter concentrates mainly on single-agent decision theory but chapters 10, 15, 16, and 18 extend the theory to environments with multiple decision makers, all of whom are concerned about model misspecification. In the process, we describe equilibrium concepts that generalize the notion of a rational expectations equilibrium to situations in which decision makers have different amounts of confidence in a common approximating model.³ Chapter

¹ Later chapters supply technical details that justify assertions made in this chapter.

 $^{^2}$ Chapter 6 studies these alternative representations in detail.

 $^{^3}$ Chapter 10 discusses competitive equilibria in representative agent economies; chapter 15 injects motives for robustness into Markov perfect equilibria for two-player dynamic games; and chapters 16 and 18 study Stackelberg and Ramsey problems. In Ramsey problems, a

^{– 17 –}

17 shows that many of the insights of this chapter extend beyond the linear quadratic setting.

2.2. Approximating models

We begin with the single agent linear quadratic problem. Let y_t be a state vector and u_t a vector of controls. A decision maker's model takes the form of a linear state transition law

$$y_{t+1} = Ay_t + Bu_t + C\check{\epsilon}_{t+1}, \tag{2.2.1}$$

where $\{\check{e}_t\}$ is an i.i.d. Gaussian vector process with mean 0 and identity contemporaneous covariance matrix. The decision maker thinks that (2.2.1) approximates another model that governs the data but that he cannot specify. How should we represent the notion that (2.2.1) is misspecified? The i.i.d. random process \check{e}_{t+1} can represent only a very limited class of approximation errors and in particular cannot depict misspecified *dynamics* such as nonlinear and time-dependent feedback of y_{t+1} on past states. To represent dynamic misspecification,⁴ we surround (2.2.1) with a set of models of the form

$$y_{t+1} = Ay_t + Bu_t + C\left(\epsilon_{t+1} + w_{t+1}\right), \qquad (2.2.2)$$

where $\{\epsilon_t\}$ is another i.i.d. Gaussian process with mean zero and identity covariance matrix and w_{t+1} is a vector process that can feed back in a possibly nonlinear way on the history of y:

$$w_{t+1} = g_t \left(y_t, y_{t-1}, \ldots \right), \tag{2.2.3}$$

government chooses among competitive equilibria of a dynamic economy. A Ramsey problem too ends up looking like a single-agent problem, the single agent being a benevolent government that faces a peculiar set of constraints that represent competitive equilibrium allocations.

 $^{^4}$ In appendix A and chapters 5 and 17, we allow a broader class of misspecifications. Chapter 17 represents the approximating model as a Markov transition density and considers misspecifications that twist probabilities over future states. When the approximating model is Gaussian, many results of this chapter survive even though (2.2.2) ignores an additional adjustment to the innovation covariance matrix of the distorted model. In many applications, the adjustment to the covariance matrix is quantitatively insignificant. It vanishes in the case of continuous time. See page 400.

where $\{g_t\}$ is a sequence of measurable functions. When (2.2.2) generates the data, it is as though the errors $\check{\epsilon}_{t+1}$ in model (2.2.1) are distributed as $\mathcal{N}(w_{t+1}, I)$ rather than as $\mathcal{N}(0, I)$. Thus, we capture the idea that the approximating model (2.2.1) is misspecified by allowing the conditional mean of the shock vector in the model (2.2.2) that actually generates the data to feed back arbitrarily on the history of the state. To express the idea that model (2.2.1) is a good approximation when (2.2.2) generates the data, we restrain the approximation errors by

$$E_0 \sum_{t=0}^{\infty} \beta^{t+1} w'_{t+1} w_{t+1} \le \eta_0, \qquad (2.2.4)$$

where E_t denotes mathematical expectation evaluated with model (2.2.2) and conditioned on $y^t = [y_t, \ldots, y_0]$. In section 2.3 and chapter 8, we shall interpret the left side of (2.2.4) as a statistical measure of the discrepancy between the distorted and approximating models.

The decision maker believes that the data are generated by a model of the form (2.2.2) with some unknown process w_t satisfying (2.2.4). The decision maker forsakes trying to improve his specification by learning because η_0 is so small that statistically it is difficult to distinguish model (2.2.2) from (2.2.1) using a time series $\{y_t\}_{t=1}^T$ of moderate size T, an idea that we develop in Chapter 8.⁵

The decision maker's distrust of his model (2.2.1) makes him want good decisions over a set of models (2.2.2) satisfying (2.2.4). Such decisions are said to be robust to model misspecification.

We compute robust decisions rules by solving one of several distinct but related two-player zero-sum games: a maximizing decision maker chooses controls $\{u_t\}$ and a minimizing player (a 'malevolent' or 'evil' agent) chooses model distortions $\{w_{t+1}\}$. The games share common players, actions, and payoffs, but assume different timing protocols. Nevertheless, as we show in chapters 6 and 7, equilibrium outcomes and decision rules for the games coincide. Important technicalities must be sorted through to verify this claim, ⁶ but their equivalence makes the games easy to solve. Computing robust decision rules comes down

⁵ However, see Hansen and Sargent (2005,nocommitXXXXX) for a way to include a robust form of learning. Also, see chapter 14.

 $^{^{6}}$ The main thing that generates this outcome is that all of them are zero-sum games, a feature that perfectly misaligns the preferences of the two players.

to solving Bellman equations for dynamic programming problems that are very similar to ones routinely used today throughout macroeconomics and applied economic dynamics. Before later chapters assemble the results needed to substantiate these claims, this chapter quickly summarizes how to compute robust decision rules with standard methods.

We begin with the ordinary linear quadratic dynamic programming problem without model misspecification, called the optimal linear regulator. Then we describe how robust decision rules can be computed by solving another optimal linear regulator problem. Next we briefly describe Lagrangian (or Hamiltonian) methods. We close by highlighting material from chapter 16 that shows how those Lagrangian methods achieve robust control of forward-looking macro models and thereby solve robust Ramsey or Stackelberg problems.

2.2.1. Dynamic programming without model misspecification

The standard dynamic programming problem assumes that the transition law is correct.⁷ Let the one-period loss function be r(y, u) = -(y'Qy + u'Ru)where the matrices Q and R are symmetric and satisfy some stabilizability and detectability assumptions set forth in chapter 3. The *optimal linear regulator* problem is

$$-y_0 P y_0 - p = \max_{\{u_t\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t r\left(y_t, u_t\right), \quad 0 < \beta < 1,$$
(2.2.5)

where the maximization is subject to (2.2.1), y_0 is given, E denotes the mathematical expectation operator evaluated with respect to the distribution of $\check{\epsilon}$, and E_0 denotes the mathematical expectation conditional on time 0 information, namely the state y_0 . Letting y^* denote next period's value of y, the linear constraints and quadratic objective function in (2.2.5), (2.2.1) imply the Bellman equation

$$-y'Py - p = \max_{u} E[r(y,u) - \beta y^{*'}Py^{*} - \beta p] | y$$
 (2.2.6)

where the maximization is subject to

$$y^* = Ay + Bu + C\check{\epsilon},\tag{2.2.7}$$

 $^{^{7}\,}$ Many technical results and computational methods for this problem are catalogued in chapter 3.

where $\check{\epsilon}$ is a random vector with mean zero and identity variance matrix.

Subject to assumptions about A, B, R, Q, β to be described in Chapter 3, some salient facts about the optimal linear regulator are these:

1. The Riccati equation. The matrix P in the value function is a fixed point of a matrix Riccati equation:

$$P = Q + \beta A' P A - \beta^2 A' P B \left(R + \beta B' P B \right)^{-1} B' P A.$$
(2.2.8)

The optimal decision rule is $u_t = -Fy_t$ where

$$F = \beta \left(R + \beta B' P B \right)^{-1} B' P A. \tag{2.2.9}$$

We can find the appropriate fixed point P and solve problem (2.2.5), (2.2.1) by iterating to convergence on the Riccati equation (2.2.8) starting from initial value $P_0 = 0$.

- 2. Certainty equivalence. In the Bellman equation (2.2.6), the scalar $p = \frac{\beta}{1-\beta}$ trace*PCC'*. The 'volatility matrix' *C* influences the value function through *p*, but not through *P*. It follows from (2.2.8), (2.2.9) that the optimal decision rule *F* is independent of the volatility matrix *C*. In (2.2.1), we have normalized by setting $E\check{\epsilon}_t\check{\epsilon}_t' = I$. Therefore, the matrix *C* determines the covariance matrix *CC'* of random shocks impinging on the system. The finding that *F* is independent of the volatility matrix *C* is known as the certainty equivalence principle: the same decision rule $u_t = -Fx_t$ emerges from stochastic ($C \neq 0$) and nonstochastic (C = 0) versions of the problem. This kind of certainty equivalence fails to describe problems that express a concern for model misspecification; but another useful kind of certainty equivalence does. See page 27.
- 3. Shadow prices. Since the value function is $-y'_0Py_0 p$, the vector of shadow prices of the initial state is $-2Py_0$. Form a Lagrangian for (2.2.1), (2.2.5) and let the vector $-2\beta^{t+1}\mu_{t+1}$ be Lagrange multipliers on the time t version of (2.2.1). First-order conditions for a saddle point of the Lagrangian can be rearranged to form a first-order vector difference equation in (y_t, μ_t) . The optimal policy solves this difference equation subject to an initial condition for y_0 and a transversality or 'detectability' condition $\sum_{t=0}^{\infty} \beta^t r(y_t, u_t) < +\infty$. On page 47 and in chapter 3, we show that subject to these boundary conditions, the vector difference equation consisting of

the first-order conditions is solved by setting $\mu_t = Py_t$, where P solves the Riccati equation (2.2.8).

2.3. Measuring model misspecification: entropy

We use entropy to measure model misspecification. To interpret our measure of entropy, we state a modified certainty equivalence principle for linear quadratic models. Although we use a statistical interpretation of entropy, by appealing to the modified certainty equivalence result to be stated on page 27, we shall be able to drop randomness from the model but still retain a measure of model misspecification that takes the form of entropy.

Let the approximating model again be (2.2.1) and let the distorted model be (2.2.2). The approximating model asserts that $w_{t+1} = 0$. For convenience, we analyze the consequences of a fixed decision rule and assume that $u_t = -Fx_t$. Let $A_o = A - BF$ and write the approximating model as

$$y_{t+1} = A_o y_t + C\check{\epsilon}_{t+1} \tag{2.3.1}$$

and a distorted model as^8

$$y_{t+1} = A_o y_t + C \left(\epsilon_{t+1} + w_{t+1} \right). \tag{2.3.2}$$

The approximating model (2.3.1) asserts that $\check{\epsilon}_{t+1} = (C'C)^{-1}C'(y_{t+1} - A_o y_t)$. When the distorted model generates the data, $y_{t+1} - A_o y_t = C\check{\epsilon}_{t+1} = C(\epsilon_{t+1} + w_{t+1})$, which implies that the disturbances under the approximating model appear to be

$$\check{\epsilon}_{t+1} = \epsilon_{t+1} + w_{t+1}, \tag{2.3.3}$$

so that misspecification manifests itself in a distortion to the conditional mean of innovations to the state evolution equation.

How close is the approximating model to the distorted model that actually governs the data? To measure the statistical discrepancy between the two models of the transition from y to y^* , we use conditional relative entropy defined as

$$I(f)(y) = \int \log\left(\frac{f(y^*|y)}{f_o(y^*|y)}\right) f(y^*|y) \, dy^*.$$

⁸ Appendix A allows a larger set of perturbations to the approximating model and gives a definition of entropy appropriate for those perturbations.

where f_o denotes the one-step transition density associated with the approximating model and f is a transition density obtained by distorting the approximating model.⁹

In the present setting, the transition density for the approximating model is

$$f_o(y^*|y) \sim \mathcal{N}(Ay + Bu, CC'),$$

while the transition density for the distorted model is¹⁰

$$f(y^*|y) \sim \mathcal{N}(Ay + Bu + Cw, CC')$$

where both u and w are measurable functions of y^t . To evaluate entropy, we first compute the ratio of probability densities (i.e., the ratio of likelihood functions) of y_{t+1} under the distorted and the approximating models, conditional on y_t . Because w_{t+1} is measurable with respect to the history y^t , then conditional on y^t , the log likelihood of y_{t+1} for the distorted model is

$$\log L^d = -\log\sqrt{2\pi} - .5\epsilon'_{t+1}\epsilon_{t+1}.$$

Using (2.3.3), the conditional log likelihood of y_{t+1} under the approximating model is

$$\log L^{a} = -\log \sqrt{2\pi} - .5 \left(\epsilon_{t+1} + w_{t+1}\right)' \left(\epsilon_{t+1} + w_{t+1}\right).$$

Therefore, the log likelihood ratio of the distorted model with respect to the approximating model is

$$\log L^{d} - \log L^{a} = .5w'_{t+1}w_{t+1} + w'_{t+1}\epsilon_{t+1}.$$
(2.3.4)

⁹ Define the likelihood ratio $m(f(y^*|y)) = \frac{f(y^*|y)}{f_0(y^*|y)}$. Then notice that

$$I(f)(y) = \int (m\log m) f_o(y^*|y) dy^* = E_{f_o}[m\log m|y],$$

where the subscript f_o means integration with respect to the approximating model f_o . Hansen and Sargent (2005a, 2005b – commit and nocommit XXXX) exploit such representations of entropy. See appendix A.

¹⁰ In a continuous time diffusion setting, Hansen, Sargent, Turmuhambetova, and Williams (2004XXX) describe how the assumption that the distorted model is difficult to distinguish statistically from the approximating model means that it can be said to be absolutely continuous over finite intervals with respect to the approximating model. They show that this implies that the perturbations must then assume a continuous time version of the form imposed here (i.e., they can alter the drift but not the volatility of the diffusion).

Define entropy $I(w_{t+1})$ as the mathematical expectation of the log likelihood ratio (2.3.4), evaluated when the data are generated by the *distorted* model. Because w_{t+1} is measurable with respect to the history y^t , averaging (2.3.4) over ϵ_{t+1} gives the expected log likelihood

$$I(w_{t+1}) = .5w'_{t+1}w_{t+1}.$$
(2.3.5)

In chapter 8, we describe how measures like (2.3.5) govern the distribution of test statistics for discriminating among models. In chapter 12, we show how the log likelihood ratio (2.3.4) also plays an important role in pricing risky securities under an approximating model when a representative agent is concerned about model misspecification.

As an intertemporal measure of the size of model misspecification, we take

$$R(w) = 2E_0 \sum_{t=0}^{\infty} \beta^{t+1} I(w_{t+1}), \qquad (2.3.6)$$

where the mathematical expectation conditioned on y_0 is evaluated with respect to the distorted model (2.3.2). Then we impose constraint (2.2.4) on the set of models or equivalently

$$R\left(w\right) \le \eta_0. \tag{2.3.7}$$

In the next section, we construct decision rules that work well over a set of models that satisfy (2.3.7). Such robust rules can be obtained by finding the best response for a maximizing player in the equilibrium of a two-player zero-sum game.
2.4. Two robust control problems

This section states two robust control problems, the constraint problem and the multiplier problem. The two problems differ in how they treat constraint (2.3.7). Under proper conditions, the two problems have identical solutions. The multiplier problem is a robust version of a stochastic optimal linear regulator. A certainty equivalence principle allows us to compute the optimal decision rule for the multiplier problem by solving a corresponding nonstochastic optimal linear regulator problem.

We state the

Constraint problem: Given an η_0 satisfying $\overline{\eta} > \eta_0 \ge 0$, a constraint problem is

$$\max_{\{u_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t r\left(y_t, u_t\right)$$
(2.4.1)

where the extremization¹¹ is subject to the distorted model (2.2.2) and the entropy constraint (2.3.7), and where E_0 , the mathematical expectation conditioned on y_0 , is evaluated with respect to the distorted model (2.2.2). Here $\bar{\eta}$ measures the largest set of perturbations against which it is possible to seek robustness.

Next we state the

Multiplier problem: Given $\theta \in (\underline{\theta}, +\infty)$, a multiplier problem is

$$\max_{\{u_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t \left\{ r\left(y_t, u_t\right) + \beta \theta w_{t+1}' w_{t+1} \right\}$$
(2.4.2)

where the extremization is subject to the distorted model (2.2.2) and the mathematical expectation is also evaluated with respect to that model.

In the max-min problem, $\theta \in (\underline{\theta}, +\infty]$ is a penalty parameter restraining the minimizing choice of the w_{t+1} sequence. The lower bound $\underline{\theta}$ is a so-called 'breakdown point' beyond which it is fruitless to seek more robustness because the minimizing agent is sufficiently unconstrained that he can push the criterion function to $-\infty$ despite the best response of the maximizing agent. Formula (7.3.13) for $\underline{\theta}$ on page 184 shows how the value of $\underline{\theta}$ depends on the return

¹¹ Following Whittle (1990), extremization means joint maximization and minimization. It is a useful term for describing saddle-point problems.

function, the discount factor, and the transition law. Tests for whether $\theta > \underline{\theta}$ are presented in formulas (6.2.9) and (6.2.10). We shall discuss the lower bound $\underline{\theta}$ and an associated upper bound $\overline{\eta}$ extensively in chapter 7.

Chapters 6 and 7 state conditions on θ and η_0 under which the two problems have identical solutions, namely, decision rules $u_t = -Fy_t$ and $w_{t+1} = Ky_t$. Chapter 6 establishes many useful facts about distinct versions of the multiplier problem that employ alternative timing protocols¹² and that justify solving the multiplier problem recursively. Let $-y'_0Py_0 - p$ be the value of problem (2.4.2). It satisfies the Bellman equation¹³

$$-y'Py - p = \max_{u} \min_{w} E\left\{r\left(y,u\right) + \theta\beta w'w - \beta y^{*'}Py^{*} - \beta p\right\}$$
(2.4.3)

where the extremization is subject to

$$y^* = Ay + Bu + C\left(\epsilon + w\right) \tag{2.4.4}$$

where * denotes next period's value, and $\epsilon \sim \mathcal{N}(0, I)$. As a tool to explore the fragility of his decision rule, in (2.4.3) the decision maker pretends that a malevolent nature chooses a feedback rule for a model-misspecification process w.

In summary, to represent the idea that model (2.2.1) is an approximation, the robust version of the linear regulator replaces the single model (2.2.1) with the set of models (2.2.2) that satisfy (2.2.4). Before describing how robust decision rules emerge from the two-player zero-sum game (2.4.2), we tell a kind of certainty equivalence that applies to the multiplier problem.

¹² For example, one timing protocol has the maximizing u player first commit at time 0 to an entire sequence, after which the minimizing w player commits to a sequence. Another timing protocol reverses the order of choices. Other timing protocols have each player choose sequentially.

¹³ In chapter 6 we show that the multiplier and constraint problems are both recursive, but that they have different state variables and different Bellman equations. Nevertheless, they lead to identical decision rules for u_t .

2.4.1. Modified certainty equivalence principle

On page 21,¹⁴ we stated a certainty equivalence principle that applies to the linear quadratic dynamic programming problem without concern for model misspecification. It fails to hold when there is concern about model misspecification. But there is another certainty equivalence principle that allows us to work with a non-stochastic version of (2.4.3), i.e., one in which $\epsilon_{t+1} \equiv 0$ in (2.4.4). In particular, it can be verified directly that precisely the same Riccati equations and the same decision rules for u_t and for w_{t+1} emerge from solving the random version of the Bellman equation (2.4.3) as would from a version that sets $\epsilon_{t+1} \equiv 0$. This fact allows us to drop ϵ_{t+1} from the state-transition equation and p from the value function -y'Py - p, without affecting formulas for the decision rules.¹⁵ Nevertheless, inspection of the Bellman equation and the formula for the decision rule for u_t show that the 'volatility matrix' C does affect the decision rule. Therefore, the version of the certainty equivalence principle stated on page 21 — that the decision rule is independent of the volatility matrix — does not hold when there are concerns about model misspecification. This is interesting because of how a desire for robustness creates an avenue for the noise statistics embedded in the volatility matrix C to impinge on decisions even with quadratic preferences and linear transition laws.¹⁶ This effect is featured in the precautionary savings model of chapter 9, a simple version of which we shall sketch in section 2.8.

 $^{^{14}}$ Tom: figure out why the pagetag is off by one page.

¹⁵ The certainty equivalence principle stated here shares with the one on page 21 the facts that P can be computed before p; it diverges from the certainty equivalence principle without robustness on page 21 in that now P and therefore F both depend on the volatility matrix C.

¹⁶ The dependence of the decision rule on the volatility matrix is an aspect that attracted researchers like Jacobson (1973) and Whittle (1990) to risk-sensitive preferences (see chapter 17).

2.5. Robust linear regulator

The modified certainty equivalence principle lets us attain robust decision rules by positing the nonstochastic law of motion

$$y_{t+1} = Ay_t + Bu_t + Cw_{t+1} \tag{2.5.1}$$

with y_0 given, where the w process is constrained by the nonstochastic counterpart to (2.2.4). By working with this nonstochastic law of motion, we obtain the robust decision rule for the stochastic problem in which (2.5.1) is replaced by (2.2.2). The approximating model assumes that $w_{t+1} \equiv 0$. Even though randomness has been eliminated, the volatility matrix C affects the robust decision rule because it influences how the specification errors w_{t+1} feed back on the state.

To induce a robust decision rule for u_t , we solve the nonstochastic version of the multiplier problem:

$$\max_{\{u_t\}} \min_{\{w_{t+1}\}} \sum_{t=0}^{\infty} \beta^t \left[r\left(y_t, u_t\right) + \theta \beta w'_{t+1} w_{t+1} \right]$$
(2.5.2)

where the extremization is subject to (2.5.1) and y_0 is given. Let $-y'_0Py_0$ be the value of (2.5.2). It satisfies the Bellman equation¹⁷

$$-y'Py = \max_{u} \min_{w} \{ r(y, u) + \theta \beta w'w - \beta y^{*'}Py^{*} \}$$
(2.5.3)

where the extremization is subject to the *distorted* model

$$y^* = Ay + Bu + Cw. (2.5.4)$$

In (2.5.3), a malevolent nature chooses a feedback rule for a model-misspecification process w. The minimization problem in (2.5.3) induces an operator $\mathcal{D}(P)$ defined by

$$-y^{*'}\mathcal{D}(P)y^{*} = \min_{w} \{\theta w'w - y^{*'}Py^{*}\}$$
(2.5.5)

where the minimization is subject to the transition law $y^* = Ay + Cw$. From the minimization problem on the right of (2.5.5), it follows that¹⁸

$$\mathcal{D}(P) = P + \theta^{-1} P C \left(I - \theta^{-1} C' P C \right)^{-1} C' P.$$
(2.5.6)

¹⁷ Notice how this is a special case of (2.4.3) with p = 0. The modified certainty equivalence principle implies that the same P matrix solves (2.5.3) and (2.4.3).

¹⁸ Before computing \mathcal{D} in formula (6.2.10), we always check whether the matrix being inverted on the right side of (2.5.6) is positive definite.

The Bellman equation (2.5.3) can then be represented as

$$-y'Py = \max_{u} \{ r(y,u) - \beta y^{*'} \mathcal{D}(P) y^{*} \}$$
(2.5.7)

where now the maximization is subject to the approximating model $y^* = Ay + Bu$ and concern for misspecification is reflected in our having replaced P with $\mathcal{D}(P)$ in the continuation value function. Notice the use of the approximating model as the transition law in the Bellman equation (2.5.7) instead of the distorted model that is used in (2.5.3), (2.5.4). The reason for the alteration in transition laws is that Bellman equation (2.5.7) encodes the activities of the minimizing agent within the operator \mathcal{D} that distorts the continuation value function.¹⁹

Define T(P) to be the operator associated with the right side of the ordinary Bellman equation (2.2.6) that we described in (2.2.8):

$$T(P) = Q + \beta A' P A - \beta^2 A' P B (R + \beta B' P B)^{-1} B' P A.$$
(2.5.8)

Then according to (2.5.7), P can be computed by iterating to convergence on the composite operator $T \circ \mathcal{D}$ and the robust decision rule can be computed by u = -Fx, where

$$F = \beta \left(R + \beta B' \mathcal{D} \left(P \right) B \right)^{-1} B' \mathcal{D} \left(P \right) A.$$
(2.5.9)

The worst case shock obeys the decision rule w = Ky, where

$$K = \theta^{-1} \left(I - \theta^{-1} C' P C \right)^{-1} C' P \left(A - BF \right).$$
 (2.5.10)

Several comments about the solution of (2.5.3) are in order.

1. Interpreting the solution. The solution of problem (2.5.2), (2.5.1) has a recursive representation in terms of a pair of feedback rules

$$u_t = -Fy_t \tag{2.5.11a}$$

$$w_{t+1} = Ky_t. (2.5.11b)$$

¹⁹ The form of (2.5.7) links this formulation of robustness to the recursive form of Jacobson's (1973) risk-sensitivity criterion proposed by Hansen and Sargent (1995), as we shall elaborate on in chapter 17.

Here $u_t = -Fy_t$ is the robust decision rule for the control u_t , while $w_{t+1} = Ky_t$ describes a worst case shock. This worst-case shock induces a distorted transition law

$$y_{t+1} = (A + CK) y_t + Bu_t. (2.5.12)$$

After having discovered (2.5.12), we can regard the decision maker as devising a robust decision rule by choosing a sequence $\{u_t\}$ to maximize

$$-\sum_{t=0}^{\infty}\beta^t \left[y_t'Qy_t + u_t'Ru_t\right]$$

subject to (2.5.12). However, as noted above, the decision maker believes that the data are actually generated by a model with an *unknown* process $w_{t+1} = \tilde{w}_{t+1} \neq 0$. It is just that by planning against the worst case process $w_{t+1} = Ky_t$, he designs a robust decision rule that performs well under a set of models. The worst-case transition law is endogenous and depends on $\theta, \beta, Q, R, A, B$. Equation (2.5.12) incorporates how the distortion wfeeds back on the state vector y; it permits w to feed back on *endogenous* components of the state, meaning that the decision maker indirectly influences future values of w through his decision rule. Allowing the distortion to depend on endogenous state variables in this way may or may not be a useful way to think about model misspecification. How useful it is depends on whether allowing w_{t+1} to feed back on endogenous components of the state vector captures plausible specifications that concern the decision maker. But there is an alternative interpretation that excludes feedback of w on endogenous state variables, which we take up next.

2. Reinterpreting the solution. We shall sometimes find it useful to reinterpret the solution of the robust linear regulator problem (2.5.1), (2.5.2) so that the decision maker believes that the distortions w do not depend on those endogenous components of the state vector whose motion his decisions affect. In particular, in chapter 6, we show that the robust decision rule $u_t = -Fy_t$ solves the ordinary linear regulator problem

$$\max_{\{u_t\}} \sum_{t=0}^{\infty} \beta^t r(y_t, u_t)$$
 (2.5.13)

subject to the distorted transition law

$$y_{t+1} = Ay_t + By_t + Cw_{t+1} (2.5.14a)$$

$$w_{t+1} = KY_t (2.5.14b)$$

$$Y_{t+1} = A^* Y_t \tag{2.5.14c}$$

where $A^* = A - BF + CK$, where (F, K) solve problem (2.5.2), (2.5.1), and where we impose the initial condition $Y_0 = y_0$. In (2.5.14), the maximizing player views Y_t as an exogenous state vector that propels the distortion w_{t+1} that twists the law of motion for state vector y_t . The solution of (2.5.13), (2.5.14) has the outcome that $Y_t = y_t \ \forall t \ge 0.^{20}$ Chapters 6 and 7 show how formulation (2.5.13), (2.5.14) emerges from a version of the multiplier problem that imposes a timing protocol in which the minimizing agent at time 0 commits to an entire sequence of distortions $\{w_{t+1}\}_{t=0}^{\infty}$ and in which it is best for the minimizing agent to make w_{t+1} obey (2.5.14b), (2.5.14c). As we shall see in chapter 7, this formulation helps us to interpret frequency domain criteria for inducing robust decision rules. In addition, the transition law (2.5.14) rationalizes a Bayesian interpretation of the robust decision maker's behavior by identifying a particular belief about the shocks for which the maximizing player's decision rule is optimal.²¹

3. *Relation to Fellner*. In the introduction to *Probability and Profit*, William Fellner wrote:

... the central problems of decision theory may be described in semiprobabilistic views. By this I mean to say that in my opinion the directly observable weights which reasonable and consistent individuals attach to specific types of prospects are not necessarily the genuine (undistorted) subjective probabilities of the prospects, although these decision weights of consistently acting individuals do bear an understandable relation to probabilities. ... the decision weights which these decision-makers attach to alternative monetary prospects need not be universally on par with probabilities attached to head-or-tails events but may in cases be

 $^{^{20}}$ In contrast to formulation (2.5.1), (2.5.2), in problem (2.5.13), (2.5.14) the maximizing agent does not believe that his decisions can influence the future position of the distortion w. Depending on the types of perturbations to the approximating model that the maximizing agent wants to protect against, we might actually prefer interpretation (2.5.1), (2.5.2) in some applications.

²¹ A decision rule is said to have a Bayesian interpretation if it is undominated in the sense of being optimal for some model. See REFERENCE XXXX (Blackwell-Girshick???).

derived from such probabilities by "slanting" or "distortion." Slanting expresses an allowance for the instability and controversial character of some types of probability judgment; the extent to which may even depend on the magnitude of the prize which is at stake when a prospect is being weighted.

Robust control theory contains concepts that embody some of Fellner's ideas. Thus, the 'decision weights' implied by the 'slanted' transition law (2.5.14) differ from the 'subjective probabilities' implied by the approximating model (2.2.1). The distortion or 'slanting' is context-specific because K depends on the parameters β, R, Q of the discounted return function.

4. Robustness bound. The minimizing player in the two player game assists the maximizing player by helping him construct a useful bound on the performance of his decision rule. Let $A_F = A - BF$ for a fixed F in a feedback rule u = -Fy. In chapter 6 on page 171, check the pagetag here: it is off a page we show that equation (2.5.7) implies that

$$-(A_Fy + Cw)' P(A_Fy + Cw) \ge -y'A'_F \mathcal{D}(P)A_Fy - \theta w'w. \quad (2.5.15)$$

The quadratic form in y on the right side is a conservative estimate of the continuation value of the state y^* under the approximating model $y^* = A_F y$.²² Inequality (2.5.15) says that the continuation value under a distorted model is at least as great as a conservative estimate of the continuation value under the approximating model, minus θ times the measure of model misspecification w'w. The parameter θ influences the conservative adjustment operator \mathcal{D} and also determines the rate at which the bound deteriorates with misspecification. Lowering θ lowers the rate at which the bound deteriorates with misspecification. Thus, (2.5.15) provides a sense in which lower values of θ provide more conservative estimates of continuation utility and therefore more robust guides to decision making.

5. Alternative games with identical outcomes. The game (2.5.2) summarized by the Bellman equation (2.5.3) is one of several two-player zero-sum games with identical lists of players, actions, and payoffs but different timing protocols. Chapter 6 describes the relationships among these games and the remarkable fact that they have identical outcomes. The analysis of chapter

²² That is, when w = 0, $-(A_F y)' \mathcal{D}(P) A_F y$ understates the continuation value.

6 justifies using recursive methods to solve all of the games. That chapter also discusses senses in which the decision maker's preferences are dynamically consistent.

6. Approximating and worst-case models. The behavior of the state under the robust decision rule and the worst case model can be represented

$$y_{t+1} = Ay_t - BFy_t + CKy_t. (2.5.16)$$

However, the decision maker does not really believe that the worst-case shock process will prevail. He uses $w_{t+1} = Ky_t$ to slant the transition law as a tool to help him construct a rule that will be robust against a range of other w_{t+1} processes that represent unknown departures from his approximating model. We want to evaluate the performance of the robust decision rule under other models. In particular, we often want to evaluate the robust decision rule when the approximating model governs the data (so that the decision maker's fears of model misspecification are actually unfounded). With the robust decision rule and the approximating model, the law of motion is

$$y_{t+1} = (A - BF) y_t. (2.5.17)$$

We obtain (2.5.17) from (2.5.16) by replacing the worst case shock Ky_t by zero. Notice that although we set K = 0 in (2.5.16) to get (2.5.17), F in (2.5.16) embodies a best response to K, and thereby reflects the agent's 'pessimistic' forecasts of future values of the state. We call (2.5.17) the approximating model under the robust decision rule and we call (2.5.16) the worst-case or distorted model under the robust decision rule.²³ In chapter 12, we use stochastic versions of both the approximating model (2.5.17) and the distorted model (2.5.16) to express alternative formulas for the prices of risky assets when consumers fear model misspecification.

7. Lower bound on θ and H_{∞} control. Starting from $\theta = +\infty$, lowering θ increases the fear of misspecification by lowering the shadow price on the norm of the control of the minimizing player. We shall see in chapter 7 that there is a lower bound for θ . This lower bound is associated with the *largest* set of alternative models, as measured by entropy, against which

²³ The model with randomness adds $C\epsilon_{t+1}$ to the right side of (2.5.17).

it is feasible to seek a robust rule: for values of θ below this bound, the minimizing agent is penalized so little that he finds it possible to choose a distortion that sends the criterion function to $-\infty$. Control theorists are interested in the cutoff value of θ because it is affiliated with a rule that is robust to the biggest allowable set of misspecifications. We describe the associated H_{∞} control theory in chapter 7. However, the applications that we are interested in usually call for values of θ that far exceed the cutoff value. We explain why in chapter 17, where we use detection error probabilities to discipline the setting for θ .

8. Risk-sensitive preferences. It is a useful fact that we can ignore doubts about model specification and instead adjust attitudes toward risk in a way that implies the decision rule and value function that come from the two player zero-sum game (2.5.2). In particular, the decision rule $u_t = -Fx_t$ that solves the robust control problem also solves a stochastic infinite horizon discounted control problem in which the decision maker has no concern about model misspecification but instead adjusts continuation values to express an additional aversion to risk. The risk adjustment is a special case of one that Epstein and Zin (1989) used to formulate their recursive specification of utility and is governed by a parameter $\sigma < 0$. If we set $\sigma = -\theta^{-1}$ from the robust control problem, we recover the same decision rule for the two problems.

The risk-sensitive decision maker trusts that the law of motion for the state is

$$y_{t+1} = Ay_t + Bu_t + C\epsilon_{t+1} \tag{2.5.18}$$

where $\{\epsilon_{t+1}\}\$ is a sequence of i.i.d. Gaussian random vectors with mean zero and identity covariance matrix. The utility index of the decision maker is defined recursively as the fixed point of recursions on

$$U_{t} = r(y_{t}, u_{t}) + \beta \mathcal{R}_{t}(U_{t+1})$$
(2.5.19)

where

$$\mathcal{R}_t \left(U_{t+1} \right) = \frac{2}{\sigma} \log E \left[\exp \left(\frac{\sigma U_{t+1}}{2} \right) \left| y^t \right]$$
(2.5.20)

and where $\sigma \leq 0$ is the risk-sensitivity parameter. When $\sigma = 0$, an application of l'hospital's rule shows that \mathcal{R}_t becomes the ordinary conditional expectation operator $E(\cdot|y^t)$. When $\sigma < 0$, \mathcal{R}_t puts an additional adjustment for risk into the assessment of continuation values.

For a quadratic r(y, u), the Bellman equation for Hansen and Sargent's (1995) risk-sensitive control problem is

$$-y'Py - \hat{p} = \max_{u} \left\{ r\left(y, u\right) + \beta \mathcal{R}\left(-y^{*'}Py^{*} - \hat{p}\right) \right\},$$
(2.5.21)

where the maximization is subject to $y^* = Ay + Bu + C\epsilon$ and ϵ is a Gaussian vector with mean zero and identity covariance matrix.

Using a result from Jacobson (1973), it can be shown that

$$\mathcal{R}(-y^{*'}Py^{*} - \hat{p}) = -(Ay + Bu)' \mathcal{D}(P)(Ay + Bu) - p(P, \hat{p}) \quad (2.5.22)$$

where \mathcal{D} is the same operator defined in (2.5.6) with $\theta = -\sigma^{-1}$, and the operator p is defined by

$$p(P, \hat{p}) = \hat{p} - \sigma^{-1} \log \det (I + \sigma C' P C).$$
 (2.5.23)

Consequently, the Bellman equation for the infinite-horizon discounted risksensitive control problem can be expressed as

$$-y'Py - \hat{p} = \max_{u} \{ r(y,u) - \beta (Ay + Bu)' \mathcal{D}(P) (Ay + Bu) - \beta p(P, \hat{p}) \}.$$
(2.5.24)

Evidently, the fixed point P satisfies $P = T \circ \mathcal{D}(P)$, and therefore is the same P that appears in the Bellman equation (2.4.3) for the robust control problem. The constant \hat{p} that solves (2.5.24) differs from the p in (2.4.3), but since they depend only on P and not on p or \hat{p} , the decision rules are the same for the two problems.

2.6. More general misspecifications

Thus far, we have permitted the decision maker to seek robustness against misspecifications that occur only as a distortion w_{t+1} to the conditional mean of the innovation to the state y_{t+1} . When the approximating model has the Gaussian form (2.2.1), this is less restrictive than may at first appear. In appendix A and chapter 17, we allow a more general class of misspecifications to the linear Gaussian model (2.2.1), but nevertheless find that important parts of the preceding results survive. For convenience, express the approximating model (2.2.1) in the compact notation

$$f_o(y^*|y) \sim \mathcal{N}(Ay + Bu, CC'),$$

which portrays the conditional distribution of next period's state as Gaussian with mean Ay + Bu and covariance matrix CC'. Let $f(y^*|y)$ be an arbitrary alternative conditional distribution that puts positive probability on the same events as the approximating model f_o . The conditional entropy of model frelative to the approximating model f_o is

$$I(f)(y) = \int \log\left(\frac{f(y^*|y)}{f_o(y^*|y)}\right) f(y^*|y) \, dy^*.$$

Entropy I(f)(y) is thus the conditional expectation of the log likelihood ratio evaluated with respect to the distorted model f. A multiplier robust control problem is associated with the following Bellman equation:

$$-y'Py - p = \max_{u} \min_{f} E\left\{r\left(y, u\right) + 2\theta\beta I\left(f\right)\left(y\right) - \beta y^{*'}Py^{*} - \beta p\right\}$$
(2.6.1)

Let $\sigma = -\theta^{-1}$ and consider the inner minimization problem, assuming that u = -Fy. On page 399, we shall show that the extremizing f is the Gaussian distribution

$$f(y^*|y) \sim \mathcal{N}\left(Ay - BFy + CKy, \hat{C}\hat{C}'\right)$$
(2.6.2)

where (F, K) are the same matrices appearing in (2.5.11),

$$\hat{C}\hat{C}' = C \left(I + \sigma C' P C\right)^{-1} C', \qquad (2.6.3)$$

and P is the same P that appears in the solution of the Bellman equation for the deterministic multiplier robust control problem (2.5.3). Equation (2.6.2) assures us that when we allow the minimizing player to choose a general misspecification $f(y^*|y)$, he chooses a Gaussian distribution with the *same* mean distortion as when we let him distort only the mean of a Gaussian conditional distribution. However, formula (2.6.3) shows that the minimizing agent would also distort the covariance matrix of the innovations, if given a chance.²⁴

The upshot of these findings is that when the conditional distribution $f(y^*|y)$ for the approximating model is Gaussian, even if we actually were to permit general misspecifications $f(y^*|y^t)$, we could compute the worst-case f by solving a deterministic multiplier robust control problem for P, F, K, and then use P to compute the appropriate adjustment to the covariance matrix (2.6.3). In chapter 12, we use some of these ideas to price assets under alternative assumptions about the set of models against which decision makers seek robustness.

2.7. A simple algorithm

Chapter 6 discusses alternative algorithms for solving (2.5.3) and relationships among them. This section describes perhaps the simplest algorithm, an adapted ordinary optimal linear regulator. Chapters 6 and 7 describe necessary technical conditions, including restrictions on the magnitude of the multiplier parameter θ .²⁵

Application of the ordinary optimal linear regulator can be justified by noting that the Riccati equation for the optimal linear regulator emerges from first-order conditions alone, and that the first-order conditions for extremizing (i.e., finding the saddle point by simultaneously minimizing with respect to wand maximizing with respect to u) the right side of (2.5.3) match those for an ordinary (non-robust) optimal linear regulator with joint control process $\{u_t, w_{t+1}\}$. This insight allows us to solve (2.5.3) by forming an appropriate optimal linear regulator.

 $^{^{24}}$ In a diffusion setting in continuous time, the minimizing agent chooses not to distort the volatility matrix because it is infinitely costly in terms of entropy. See Hansen, Sargent, Turmuhambetova, and Williams (2005) and Anderson, Hansen, and Sargent (2003).

²⁵ The Matlab program olrprobust.m described in the appendix implements this algorithm; doublex9.m implements a doubling algorithm of the kind described in chapter 3 and Hansen and Sargent (XXXXbook); please note that doublex9.m solves a *minimum* problem and that $-\theta^{-1} \equiv \sigma < 0$ connotes a fear of model misspecification.

Thus, put the Bellman equation (2.5.3) into a more compact form by defining

$$\tilde{B} = \begin{bmatrix} B & C \end{bmatrix} \tag{2.7.1a}$$

$$\tilde{R} = \begin{bmatrix} R & 0\\ 0 & -\beta\theta I \end{bmatrix}$$
(2.7.1*b*)

$$\tilde{u}_t = \begin{bmatrix} u_t \\ w_{t+1} \end{bmatrix}. \tag{2.7.1c}$$

Let ext denote 'extremization' – maximization with respect to u, minimization with respect to w. The Bellman equation can be written

$$-y'Py = \operatorname{ext}_{\tilde{u}}\left\{-y'Qy - \tilde{u}\tilde{R}\tilde{u} - \beta y^{*'}Py^{*}\right\}$$
(2.7.2)

where the extremization is subject to

$$y^* = Ay + \tilde{B}\tilde{u}. \tag{2.7.3}$$

The first-order conditions for problem (2.7.2), (2.7.3) imply the matrix Riccati equation

$$P = Q + \beta A' P A - \beta^2 A' P \tilde{B} \left(\tilde{R} + \beta \tilde{B}' P \tilde{B} \right)^{-1} \tilde{B}' P A \qquad (2.7.4)$$

and the formula for \tilde{F} in the decision rule $\tilde{u}_t = -\tilde{F} y_t$

$$\tilde{F} = \beta \left(\tilde{R} + \beta \tilde{B}' P \tilde{B} \right)^{-1} \tilde{B}' P A.$$
(2.7.5)

Partitioning \tilde{F} , we have

$$u_t = -Fy_t \tag{2.7.6a}$$

$$w_{t+1} = Ky_t. (2.7.6b)$$

The decision rule $u_t = -Fy_t$ is the robust rule. As mentioned above, $w_{t+1} = Ky_t$ provides the θ -constrained worst-case specification error. We can solve the Bellman equation by iterating to convergence on the Riccati equation (2.7.4), or by using one of the faster computational methods described in chapter 3.

2.7.1. Interpretation of the simple algorithm

The adjusted Riccati equation (2.7.4) is an augmented version of the Riccati equation (2.2.8) that is associated with the ordinary optimal linear regulator. The right side of equation (2.7.4) defines one step on the composite operator $T \circ \mathcal{D}$ where T and \mathcal{D} are defined in (2.5.8) and (2.5.5).²⁶ Chapter 7 connects the \mathcal{D} operator to Hansen and Sargent's (1995) discounted version of the risk-sensitive preferences of Jacobson (1973) and Whittle (1990).

2.8. Example: robustness and discounting in a permanent income model

This section illustrates various aspects of robust control theory in the context of a linear-quadratic version of a simple permanent income model.²⁷ In the basic permanent income model, a consumer applies a single marginal propensity to consume to the sum of his financial wealth and his human wealth, where human wealth is defined as the expected present value of his labor (or endowment) income discounted at the same risk-free rate of return that he earns on his financial assets. In the usual permanent income model without a concern about robustness, the consumer has no doubts about the probability model used to form the conditional expectation of discounted future labor income. Instead, we assume that the consumer doubts that model and therefore forms forecasts of future income by using a conditional probability distribution that is twisted or slanted relative to his approximating model for his endowment. Otherwise, the consumer behaves as an ordinary permanent income consumer.

His slanting of conditional probabilities leads the consumer to engage in a form of precautionary savings that under the approximating model for his endowment process tilts his consumption profile toward the future relative to what it would be without a concern about misspecification of that process. Indeed, so far as his consumption and saving program is concerned, activating a concern about robustness is equivalent with making the consumer more patient.

 $^{^{26}\,}$ This can be verified by unstacking the matrices in (2.7.4). See page 170 in chapter 6.

²⁷ See Sargent (1987) and Hansen, Roberds, and Sargent (1991) for accounts of the connection between the permanent income consumer and Barro's (1979) model of tax smoothing. See Aiyagari, Marcet, Sargent, and Seppälä (2002) for a deeper exploration of the connections.

However, that is not the end of the story. Chapter 12 shows that attributing a concern about robustness to a representative consumer has different effects on asset prices than are associated with varying his discount factor.

2.8.1. The LQ permanent income model

In Hall's (1978) linear-quadratic permanent income model, a consumer receives an exogenous endowment $\{d_t\}$ and wants to allocate it between consumption c_t and savings k_t to maximize

$$-E_0 \sum_{t=0}^{\infty} \beta^t \left(c_t - b \right)^2, \beta \in (0,1).$$
(2.8.1)

We simplify the problem by assuming that the endowment is a first-order autoregression. Thus, the household faces the state transition laws

$$k_t + c_t = Rk_{t-1} + d_t \tag{2.8.2a}$$

$$d_{t+1} = \mu_d \left(1 - \rho \right) + \rho d_t + c_d \left(\epsilon_{t+1} + w_{t+1} \right), \qquad (2.8.2b)$$

where R > 1 is a time-invariant gross rate of return on financial assets k_{t-1} held at the end of period t-1, and $|\rho| < 1$ describes the persistence of his endowment. In (2.8.2b), w_{t+1} is a distortion to the mean of the endowment that represents possible model misspecification. We use $\sigma = -\theta^{-1}$ to parameterize the consumer's desire for robustness. Soon we'll confirm how easily this problem maps into the robust linear regulator. But first we'll use classical methods to elicit some useful properties of the consumer's decisions when $\sigma = 0$.

2.8.2. Solution when $\sigma = 0$

We first solve the household's problem *without* a concern about robustness by setting $\theta^{-1} \equiv \sigma = 0$. Define the marginal utility of consumption as $\mu_{ct} = b - c_t$. The household's Euler equation is

$$E_t \mu_{c,t+1} = (\beta R)^{-1} \mu_{ct}, \qquad (2.8.3)$$

where E_t is the mathematical expectation operator conditioned on date t information. Treating (2.8.2a) as a difference equation in k_t , solving it forward in time, and taking conditional expectations on both sides gives

$$k_{t-1} = \sum_{j=0}^{\infty} R^{-(j+1)} E_t \left(c_{t+j} - d_{t+j} \right).$$
(2.8.4)

Solving (2.8.3) and (2.8.4) and using $\mu_{ct} = b - c_t$ implies

$$\mu_{ct} = -\left(1 - R^{-2}\beta^{-1}\right) \left(Rk_{t-1} + E_t \sum_{j=0}^{\infty} R^{-j} \left(d_{t+j} - b\right)\right).$$
(2.8.5)

Equations (2.8.3) and (2.8.5) can be used to deduce the following representation for μ_{ct}

$$\mu_{c,t+1} = (\beta R)^{-1} \mu_{c,t} + \nu \epsilon_{t+1}.$$
(2.8.6)

We provide a formula for the scalar ν in (2.8.11) below.

Given an initial condition $\mu_{c,0}$, equation (2.8.6) describes the consumer's optimal behavior; $\mu_{c,0}$ can be determined by solving (2.8.5) at t = 0. It is easy to use (2.8.5) to deduce an optimal consumption rule of the form

$$c_t = gy_t$$

where g is a vector conformable to the pertinent state vector y. In the case $\beta R = 1$ that was analyzed by Hall (1978), (2.8.6) implies that the marginal utility of consumption μ_{ct} is a martingale under the approximating model, which because $\mu_{ct} = b - c_t$ in turn implies that consumption itself is a martingale.

2.8.3. Linear regulator for permanent income model

This problem is readily mapped into a linear regulator in which the marginal utility of consumption $b - c_t$ is the control. Express the transition law for k_t as

$$k_t = Rk_{t-1} + d_t - b + (c_t - b).$$

Define the state as $y'_t = \begin{bmatrix} 1 & k_{t-1} & d_t \end{bmatrix}'$ and the control as $u_t = \mu_{ct} \equiv (b - c_t)$ and express the state transition law as $y_{t+1} = Ay_t + Bu_t + C(\epsilon_{t+1} + w_{t+1})$ or

$$\begin{bmatrix} 1\\k_t\\d_{t+1} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0\\-b & R & 1\\0 & 0 & \rho \end{bmatrix} \begin{bmatrix} 1\\k_{t-1}\\d_t \end{bmatrix} + \begin{bmatrix} 0\\1\\0 \end{bmatrix} (b-c_t) + \begin{bmatrix} 0\\0\\c_d \end{bmatrix} (\epsilon_{t+1} + w_{t+1})$$
(2.8.7)

This equation defines the triple (A, B, C) associated with a robust linear regulator. For the objective function, (2.8.1) implies that we should let r(y, u) = -y'Ry - u'Qu where $R = 0_{3\times3}$ and Q = 1.

We can obtain a robust rule by using the robust linear regulator and setting $\sigma < 0$. The solution of the robust linear regulator problem is a linear decision rule for the control μ_{ct} :

$$\mu_{ct} = -Fy_t. \tag{2.8.8}$$

Under the approximating model, the law of motion of the state is then

$$y_{t+1} = (A - BF) y_t + C\epsilon_{t+1}.$$
(2.8.9)

Equations (2.8.8) and (2.8.9) imply that

$$\mu_{c,t+1} = -F \left(A - BF \right) y_t - FC \epsilon_{t+1}. \tag{2.8.10}$$

Comparing (2.8.10) and (2.8.6) shows that $-F(A - BF) = -(\beta R)^{-1}F$ and

$$\nu = -FC, \tag{2.8.11}$$

which is the promised formula for ν .

2.8.4. Effects on consumption of concern about misspecification

To understand the effects on consumption of a concern about robustness, we use as a benchmark Hall's assumption that $\beta R = 1$ and no concern about robustness ($\sigma = 0$). In that case, the multiplier μ_{ct} and consumption c_t are both driftless random walks. To be concrete, we set parameters to be consistent with ones calibrated from post-World War II U.S. time series by Hansen, Sargent, and Tallarini (1999) for a more general permanent income model. HST set $\beta = .9971$ and fit a two-factor model for the endowment process; each factor is a second order autoregression. To simplify that specification, we replace this estimated two-factor endowment process with the population first-order autoregression one would obtain if that two factor model actually generated the data. That is, we use the population moments implied by Hansen, Sargent, and Tallarini's (HST's) estimated endowment process to fit the first-order autoregressive process (2.8.2b) with $w_{t+1} \equiv 0$. Ignoring constant terms, we obtain the endowment process $d_{t+1} = .9992d_t + 5.5819\epsilon_{t+1}$ where ϵ_{t+1} is an i.i.d. scalar process with mean zero and unit variance.²⁸ We use $\hat{\beta}$ to denote HST's value of $\beta = .9971$. Throughout, we suppose that $R = \hat{\beta}^{-1}$.

We now consider three cases.

• The $\beta R = 1, \sigma = 0$ case studied by Hall (1978). With $\beta = \hat{\beta}$, we compute that the marginal utility of consumption follows the law of motion

$$\mu_{c,t+1} = \mu_{c,t} + 4.3825\epsilon_{t+1} \tag{2.8.12}$$

where we compute the coefficient 4.3825 on ϵ_{t+1} by noting that it equals -FC by formula (2.8.11).

• A version of Hall's $\beta R = 1$ specification with a concern about misspecification. Retaining $\hat{\beta}R = 1$, we activate a concern about robustness by setting $\sigma = \hat{\sigma} - 2E - 7 < 0.^{29}$ We now compute that ³⁰

$$\mu_{c,t+1} = .9976\mu_{c,t} + 8.0473\epsilon_{t+1}. \tag{2.8.13}$$

ŀ

²⁸ We computed ρ , c_d by calculating autocovariances implied by HST's specification, then used them to calculate the implied population first-order autoregressive representation.

 $^{^{29}\,}$ We discuss how to calibrate $\,\sigma\,$ in chapters 9, 12, and 17.

³⁰ We can confirm this formula computationally as follows. Use doublex9 to solve the robust optimal linear regulator and compute representations $\mu_{c,t} = -Fy_t$ and compare it to the term $F(A - BF)y_t$ on the right side of (2.8.10) to discover that F(A - BF) = .9976F i.e., the coefficients are proportional with .9976 being the factor of proportionality.

When $b - c_t > 0$, this equation implies that $E_t(b - c_{t+1}) = .9976(b - c_t) < (b - c_t)$ which in turn implies that $E_t c_{t+1} > c_t$. Thus, the effect of activating a concern about robustness is to put upward drift into the consumption profile, a manifestation of a type of 'precautionary savings' that comes from the consumer's fear of misspecification of the endowment process.

• A case that raises the discount factor relative to the $\beta R = 1$ benchmark prevailing in Hall's model, but withholds a concern about robustness. In particular, while we set $\sigma = 0$ we increase β to $\tilde{\beta} = .9995$. Remarkably, with $(\sigma, \beta) = (0, \tilde{\beta})$, we compute that $\mu_{c,t+1}$ obeys exactly (2.8.13).³¹ Thus, starting from $(\sigma, \beta) = (0, \hat{\beta})$, in so far as the effects on consumption and saving are concerned, activating a concern about robustness by lowering σ while keeping β constant is evidently equivalent to keeping $\sigma = 0$ but *increasing* the discount factor to a particular $\tilde{\beta} > \hat{\beta}$.

These numerical examples illustrate what is true more generally, that in the permanent income model an increased concern about robustness has effects on (c_t, k_{t+1}) that operate exactly like an increase in the discount factor β . In chapter 9, we extend these numerical examples analytically within a broader class of permanent income models. In particular, let $\alpha^2 = \nu'\nu$ and suppose that instead of the particular pair $(\hat{\sigma}, \hat{\beta})$, where $(\hat{\sigma} < 0)$, we use the pair $(0, \tilde{\beta})$, where $\tilde{\beta}$ satisfies:

$$\tilde{\beta}(\sigma) = \frac{\hat{\beta}\left(1+\hat{\beta}\right)}{2\left(1+\sigma\alpha^2\right)} \left[1 + \sqrt{1-4\hat{\beta}\frac{1+\sigma\alpha^2}{\left(1+\hat{\beta}\right)^2}}\right].$$
(2.8.14)

Then the laws of motion for $\mu_{c,t}$, and therefore the decision rules for c_t , are identical across these two specifications about concerns about robustness. We establish formula (2.8.14) in appendix B of chapter 9.

³¹ We discover this computationally using the method of the previous footnote.

2.8.5. Equivalence of quantities but not continuation values

We have seen that, holding other parameters constant, there exists a locus of (σ,β) pairs that imply the same consumption, saving programs. It can be verified that the P matrices appearing in the quadratic forms in the value function are identical for the $(\hat{\sigma}, \hat{\beta})$ and $(0, \tilde{\beta})$ problems. However, in terms of their implications for pricing claims on risky future payoffs, it is significant that the $\mathcal{D}(P)$ matrices differ across such (σ, β) pairs. For the $(0, \tilde{\beta})$ pair, $P = \mathcal{D}(P)$. However, when $\sigma < 0$, $\mathcal{D}(P)$ differs from P. As we shall see in chapter 12, when we interpret (2.8.1), (2.8.2) as a planning problem, $\mathcal{D}(P)$ encodes the shadow prices that can be converted into competitive equilibrium state-date prices that can then be used to price uncertain claims on future consumption. Thus, although the $(\hat{\sigma}, \hat{\beta})$ and $(0, \tilde{\beta})$ parameter settings imply identical savings and consumption plans, they imply different valuations of risky future consumption payoffs. In chapter 9, we use this fact to study how a concern about robustness influences the theoretical value of the market price of macroeconomic risk and the equity premium.

2.8.6. Distorted endowment process

On page 30, we described a particular distorted transition law associated with the worst case shocks $w_{t+1} = Ky_t$. If the decision maker solves an ordinary dynamic programming program without a concern about misspecification but substitutes the distorted transition law for the one given by his approximating model, he attains a robust decision rule. Thus, when $\sigma < 0$, instead of facing the transition law (2.8.7) that prevails under the approximating model, the household would use the distorted transition law²⁹

$$\begin{bmatrix} y_{t+1} \\ Y_{t+1} \end{bmatrix} = \begin{bmatrix} A & CK \\ 0 & (A - BF + CK) \end{bmatrix} \begin{bmatrix} y_t \\ Y_t \end{bmatrix} + \begin{bmatrix} B \\ 0 \end{bmatrix} \mu_{ct} + \begin{bmatrix} C \\ C \end{bmatrix} \epsilon_{t+1}.$$
 (2.8.15)

For our numerical example with $\sigma = -2E - 7$, we would have $A - BF + CK = \begin{bmatrix} 1.0000 & 0 & 0\\ 15.0528 & 0.9976 & -0.4417\\ -0.0558 & 0.0000 & 1.0016 \end{bmatrix}$ and $CK = \begin{bmatrix} 0 & 0 & 0\\ 0 & 0 & 0\\ -0.0558 & 0.0000 & 0.0024 \end{bmatrix}$. Notice the pattern of zeros in CK, which shows that the distortion to the law of

 $^{^{29}}$ This is not a minimal state representation because we have not eliminated the constant from the Y component of the state.

motion of the state affects only the d_t component of the component of the state y. The components Y of the state are information variables that account for the dynamics in the misspecification imputed by the worst case shock w. In chapter 9, we shall analyze the behavior of the endowment process under the distorted model (2.8.15).

It is useful to consider our observational equivalence result in light of the distorted law of motion (2.8.15). Let \hat{E}_t denote a conditional expectation with respect to the distorted transition law (2.8.15) for the endowment shock and let E_t denote the expectation with respect to the approximating model. Then the observational equivalence of the pairs $(\hat{\sigma}, \hat{\beta})$ and $(0, \tilde{\beta})$ means that the following two versions of (2.8.5) imply the same μ_{ct} processes:

$$\mu_{ct} = -\left(1 - R^{-2}\hat{\beta}^{-1}\right) \left(Rk_{t-1} + \hat{E}_t \sum_{j=0}^{\infty} R^{-j} \left(d_{t+j} - b\right)\right)$$

and

$$\mu_{ct} = -\left(1 - R^{-2}\tilde{\beta}^{-1}\right) \left(Rk_{t-1} + E_t \sum_{j=0}^{\infty} R^{-j} \left(d_{t+j} - b\right)\right).$$

For both of these expressions to be true, the effect on \hat{E} of setting σ less than zero must be offset by the effect of raising β from $\hat{\beta}$ to $\tilde{\beta}$.

2.8.7. Representing misspecification: a Stackelberg formulation

In chapters 6 and 7, we show the equivalence of outcomes under different timing protocols for the two-player zero-sum games. In appendix B of chapter 9, we shall use a Stackelberg game to establish the observational equivalence for consumption, savings plans of $(0, \tilde{\beta})$ and $(\hat{\sigma}, \hat{\beta})$ pairs. The minimizing player's problem in the Stackelberg game can be represented as

$$\min_{\{w_{t+1}\}} - \sum_{t=0}^{\infty} \hat{\beta}^t \left\{ \mu_{ct}^2 + \hat{\beta} \sigma^{-1} w_{t+1}^2 \right\}$$
(2.8.16)

subject to

$$\mu_{c,t+1} = \left(\tilde{\beta}R\right)^{-1} \mu_{c,t} + \nu w_{t+1}.$$
(2.8.17)

Equation (2.8.17) is the consumption Euler equation of the maximizing player. Under the Stackelberg timing, the minimizing player commits to a sequence $\{w_{t+1}\}_{t=0}^{\infty}$ that the maximizing player takes as given. The minimizing player determines that sequence by solving (2.8.16), (2.8.17). The worst case shock that emerges from this problem satisfies $w_{t+1} = k\mu_{ct}$ and is identical to the worst case shock $w_{t+1} = Ky_t$ that emerges from the robust linear regulator for the consumption problem.

2.9. Stabilizing property of shadow price Py_t

In chapter 3, we solve problem (2.5.1), (2.5.2) with a Lagrangian method that provides a fast way to compute P and gives insights about a recursive representation $\mu_t = Py_t$, where $-2\beta^{t+1}\mu_{t+1}$ is the vector of shadow prices on the time t+1 state vector. The Lagrangian formulation is also convenient for designing decision rules for Ramsey and Stackelberg problems, as we shall show in section 2.10 and chapter 16. Form the Lagrangian

$$\mathcal{L} = -\sum_{t=0}^{\infty} \beta^t \left[y_t' Q y_t + u_t' R u_t + 2\beta \mu_{t+1}' \left(A y_t + B u_t + C w_{t+1} - y_{t+1} \right) - \theta \beta w_{t+1}' w_{t+1} \right].$$
(2.9.1)

We want to maximize (2.9.1) with respect to sequences for u_t and y_{t+1} and minimize it with respect to a sequence for w_{t+1} . The first-order conditions with respect to u_t, y_t, w_{t+1} , respectively, are:

$$0 = Ru_t + \beta B' \mu_{t+1} \tag{2.9.2a}$$

$$\mu_t = Qy_t + \beta A' \mu_{t+1} \tag{2.9.2b}$$

$$0 = \beta \theta w_{t+1} - \beta C' \mu_{t+1}. \tag{2.9.2c}$$

Solving (2.9.2*a*) and (2.9.2*c*) for u_t and w_{t+1} and substituting into (2.5.1) gives

$$y_{t+1} = Ay_t - \beta \left(BR^{-1}B' - \beta^{-1}\theta^{-1}CC' \right) \mu_{t+1}.$$
 (2.9.3)

Write (2.9.3) as

$$y_{t+1} = Ay_t - \beta \tilde{B} \tilde{R}^{-1} \tilde{B}' \mu_{t+1}.$$
 (2.9.4)

We represent the system formed by (2.9.2b) and (2.9.4) as

$$\begin{bmatrix} I & \beta B \dot{R}^{-1} \dot{B}' \\ 0 & \beta A' \end{bmatrix} \begin{bmatrix} y_{t+1} \\ \mu_{t+1} \end{bmatrix} = \begin{bmatrix} A & 0 \\ -Q & I \end{bmatrix} \begin{bmatrix} y_t \\ \mu_t \end{bmatrix}$$
(2.9.5)

 or

48

$$L^* \begin{bmatrix} y_{t+1} \\ \mu_{t+1} \end{bmatrix} = N \begin{bmatrix} y_t \\ \mu_t \end{bmatrix}.$$
 (2.9.6)

We want to find a 'stabilizing' solution of (2.9.6), i.e., one that satisfies³⁰

$$\sum_{t=0}^\infty \beta^t y_t' y_t < +\infty.$$

Chapter 3 shows that the stabilizing solution satisfies $\mu_t = Py_t$, where P solves the matrix Riccati equation (2.7.4). Briefly, the generalized eigenvalues of (L^*, N) occur in $\sqrt{\beta}$ -symmetric pairs (i.e., $(\lambda_i, \lambda_{-i})$ such that if λ_i is an eigenvalue, another eigenvalue is $\lambda_{-i} = \frac{1}{\beta\lambda_i}$). The stabilizing solution solves stable roots backward and unstable roots forward by imposing initial conditions satisfying $\mu_0 = Py_0$. This condition replicates itself over time in the sense that

$$\mu_t = P y_t, \tag{2.9.7}$$

and implies that $\sum_{t=0}^{\infty} \beta^t y'_t y_t < \infty$.

In summary, the solution of the nonstochastic multiplier problem is given by the feedback rule

$$\begin{bmatrix} u_t \\ w_{t+1} \end{bmatrix} = -\tilde{F}y_t \tag{2.9.8}$$

where \tilde{F} depends on P through (2.7.5). We can find P either by solving a Riccati equation or by using a method that rearranges the generalized eigenvectors of L^*, N .

 $[\]overline{30}$ Chapter 3 describes the detectability and stabilizability conditions that make this restriction equivalent with $\sum_{t=0}^{\infty} \beta^t r(y_t, u_t) < +\infty$.

2.10. Forward looking models

The basic robust control problem with Bellman equation (2.5.3) pertains to a single decision maker. For macroeconomic applications with a representative agent in an economy without distortions, (2.5.3) can be used to compute equilibrium allocations and prices (for elaboration and examples see chapters 9 and 12 as well as Hansen, Sargent, and Tallarini (1999)). However, even with a representative agent, to analyze so-called Ramsey problems where there are distortions, say flat rate taxes on labor or capital, (2.5.3) must be modified. For a Ramsey problem, the robust decision maker is a government that wants to devise a plan to which it commits at time 0, taking into account the 'forward looking' behavior of private agents whose behavior is summarized by Euler equations that include the government's policy instruments u_t as 'forcing variables'. In chapter 16, we describe how to solve such robust policy design problems. We formulate the government's problem as a Lagrangian and note how the private sector's forward-looking behavior formally transforms some of the state variables in an optimal linear regulator into 'jump' variables, while converting some Lagrange multipliers into 'state variables.' Chapter 16 reviews the interesting intellectual history of the Lagrangian formulation for such problems on both sides of the Atlantic. In this section, we explain the basic idea, whose implication is that robust Ramsey policies can be computed easily by solving and appropriately manipulating an associated ordinary optimal linear regulator problem.

Here is the basic idea. In a forward looking model, we can partition the state $y = \begin{bmatrix} z \\ x \end{bmatrix}$. The z variables are true state variables, being inherited from the past, but the n_x variables x are 'jump variables' that adjust to clear markets at t, e.g., prices and quantities. The last n_x equations of (2.5.1) include descriptions of the forward-looking behavior of the private sector, e.g., private agents' Euler equations.³¹

We need n_x additional state variables. To get them, we look to the last n_x Lagrange multipliers in (2.9.1), which we call μ_{xt} , that adhere to 'implementability constraints' that the private sector's Euler equations impose on the

 $^{^{31}}$ When private agents also have concerns about robustness, some of these Euler equations pertain to *their* worst case shocks, which in general differ from the worst case shocks of the Stackelberg leader. See chapter 16.

Ramsey plan.³² The implementability constraints are in effect promise keeping constraints that require the government to confirm the private sector's past expectations about the government's setting of the current value of its policy instrument, expectations that were incorporated into past decisions of the private sector. The implementability multipliers μ_{xt} are the missing state variables. These multipliers on the promise keeping constraints encode the effects on private agents' past decisions of government promises about future policies.

Let $\mu_t = \begin{bmatrix} \mu_{zt} \\ \mu_{xt} \end{bmatrix}$. Here $-2\beta^{t+1}\mu_{zt}$ are shadow prices on the true state variables at t + 1 and $-2\beta^{t+1}\mu_{xt}$ are shadow prices on the jump variables at time t + 1, being the 'implementability multipliers'. The Ramsey problem can be written in the form (2.5.2), (2.5.1). The first-order conditions continue to be (2.9.6) and the solution requires that (y_t, μ_t) satisfy (2.9.7), where P still solves the Riccati equation associated with the Bellman equation (2.7.2). At this point the procedure for solving the robust Ramsey problem departs from that for the linear regulator. We must use (2.9.6) to solve for the jump variable x. With this purpose, write the last n_x equations of (2.9.6) as

$$\mu_{xt} = P_{21}z_t + P_{22}x_t$$

or

$$x_t = -P_{22}^{-1}P_{21}z_t + P_{22}^{-1}\mu_{xt}.$$
 (2.10.1)

Using (2.10.1), the solution of the robust Ramsey problem is

$$\begin{bmatrix} z_{t+1} \\ \mu_{x,t+1} \end{bmatrix} = \begin{bmatrix} I & 0 \\ P_{21} & P_{22} \end{bmatrix} A_o \begin{bmatrix} I & 0 \\ -P_{22}^{-1}P_{21} & P_{22}^{-1} \end{bmatrix} \begin{bmatrix} z_t \\ \mu_{xt} \end{bmatrix}$$
(2.10.2a)

$$x_t = \begin{bmatrix} -P_{22}^{-1}P_{22} & P_{22}^{-1} \end{bmatrix} \begin{bmatrix} z_t \\ \mu_{xt} \end{bmatrix}$$
(2.10.2b)

where $A_o = (A - B\tilde{F}_1 - C\tilde{F}_2)$ for the distorted or worst-case model and $A_o = (A - B\tilde{F}_1)$ for the robust rule under the approximating model.

The decision rule and worst case model distortion w_{t+1} for the Ramsey planner (a.k.a. the Stackelberg leader) can be represented

$$\begin{bmatrix} u_t \\ w_{t+1} \end{bmatrix} = \begin{bmatrix} -\tilde{F}_1 \\ -\tilde{F}_2 \end{bmatrix} \begin{bmatrix} I & 0 \\ -P_{22}^{-1}P_{21} & P_{22}^{-1} \end{bmatrix} \begin{bmatrix} z_t \\ \mu_{xt} \end{bmatrix}.$$
 (2.10.3)

 $^{^{32}}$ In chapter 16, we show how these Euler equations need to be adjusted when the competitive representative agent also has concerns about model misspecification.

Chapter 16 shows that by eliminating μ_{xt} , the robust decision rule can be represented in the form

$$u_t = \rho u_{t-1} + \alpha_0 z_t + \alpha_1 z_{t-1}.$$

Here the history dependence of the decision rule is captured through the dependence on the lagged instrument u_{t-1} . Chapter 16 gives an alternative representation for the worst-case shock

$$w_t = \nu u_{t-1} + \gamma_0 z_t + \gamma_1 z_{t-1}.$$

2.11. Concluding remarks

The discounted dynamic programming problem for quadratic returns and a linear transition function is called the optimal linear regulator problem. This problem is widely used throughout macroeconomics and applied dynamics. For linear-quadratic problems, robust decision rules can be constructed by thoughtfully using the optimal linear regulator. This is true both for single-agent problems and for some Ramsey and Stackelberg problems. The optimal linear regulator has other uses too. In chapters 4, 13, and 14 we describe filtering problems. Via the concept of duality explained there, the linear regulator can also be used to solve such filtering problems, including ones with a desire for estimates that are robust to model misspecification.

A. Link to a stochastic formulation

Although this chapter has focused primarily on nonstochastic formulations of a linear-quadratic robust control theory, many of our results also pertain to a stochastic version of our problem in which i.i.d. Gaussian shocks also occur in the transition equation for the approximating model. In this stochastic setting, there is a natural way to formulate perturbations to the controlled stochastic process implied by the approximating model and a way to measure those perturbations in terms of entropy. In a linear-quadratic-Gaussian setting, the formulas obtained from solving the nonstochastic problem contain all of the information needed to solve the corresponding stochastic problem. This appendix outlines these connections. For a more complete account, see Hansen and Sargent (2005a, 2005b XXXX).

The approximating model (2.5.18) and the y^{t-1} -measurable control process $\{u^t\}$ induces a probability measure over the infinite history y^{∞} . Hansen and Sargent (2005) show how to represent distortions of the approximating model in terms of positive martingales $\{M_t\}$ defined on the same probability space as for the approximating model. The martingale has the recursive representation

$$M_{t+1} = m_{t+1}M_t,$$

where m_{t+1} is a y^{t+1} measurable random variable with $Em_{t+1}|y^t = 1$ and $M_0 = 1$. We take M_t as a likelihood ratio for distorting the probability of the joint distribution over y^t conditioned on y_0 that is implied by the approximating model. Define the entropy of the distortion associated with M_t as the expected log-likelihood ratio with respect to the distorted distribution, which can be expressed as $E(M_t \log M_t | y_0)$. Corresponding to the factorization $M_t = M_0 \prod_{j=1}^t m_j$ is the following decomposition of entropy:

$$E(M_t \log M_t | y_0) = \sum_{j=0}^{t-1} E\left[M_j E\left(m_{j+1} \log m_{j+1} | y^j\right) | y_0\right]$$

Here $E[m_{t+1} \log m_{t+1} | y^t]$ is the conditional relative entropy of a perturbation to the one-step transition density associated with the approximating model. Notice the absence of discounting on the right side. To get a recursive formulation of stochastic robust control that sustains an enduring concern about model misspecification, Hansen and Sargent (2005) advocate using a discounted version of the object on the right side to penalize a malevolent player's choice of a sequence of increments $\{m_{t+1}\}$. This leads us to pose a two-player zero-sum game of the form

$$\min_{\{m_{t+1}\}} \max_{\{u_t\}} \sum_{t=0}^{\infty} \sum_{y^t} \beta^t M_t \left\{ r\left(y_t, u_t\right) + \theta \beta m_{t+1} \log m_{t+1} \right\} \pi \left(y^t | u^{t-1}\right) \quad (2.A.1)$$

subject to

$$M_{t+1} = m_{t+1}M_t, (2.A.2)$$

where $Em_{t+1}|y^t = 1$, $\pi(y^t|u^{t-1})$ is the probability over the history $y^t = y_t, \ldots, y_0$ induced by u^{t-1} and the law of motion

$$y_{t+1} = Ay_t + Bu_t + C\epsilon_{t+1}, (2.A.3)$$

where ϵ_{t+1} is an i.i.d. Gaussian vector with mean zero and identity covariance matrix. Guess that the value function has the multiplicative form $W(M_0, y_0) = M_0 V(y_0)$ and express the problem as

$$M_0 V(y_0) = \min_{\{m_{t+1}\}} \max_{\{u_t\}} E_0 \sum_{t=0}^{\infty} \beta^t M_t \{ r(y_t, u_t) + \theta \beta m_{t+1} \log m_{t+1} \}$$
(2.A.4)

subject to (2.A.2) and (2.A.3). The Bellman equation for this problem is

$$M_{t}V(y_{t}) = \min_{m_{t+1}} \max_{u_{t}} M_{t} \left\{ r(y_{t}, u_{t}) + \beta E\left[m_{t+1}V(y_{t+1}) + \theta m_{t+1}\log m_{t+1}\right] \right\}$$
(2.A.5)

where $V(y) = -y'Py - \rho$ and the extremization is subject to (2.A.2) and (2.A.3). Since M_t enters linearly, the Bellman equation can be rewritten

$$V(y_t) = \min_{m_{t+1}} \max_{u_t} \left\{ r(y_t, u_t) + \beta E[m_{t+1}V(y_{t+1}) + \theta m_{t+1}\log m_{t+1}] \right\} \quad (2.A.6)$$

and the minimizing increment m_{t+1} will depend on y_t but not M_t . Let w_{t+1} and Σ_t obey formulas (2.A.10) and (2.A.11) below. The Gaussian conditional distribution and the quadratic return function imply that the minimizing m_{t+1}

$$is^{33}$$

$$m_{t+1}^* \propto \exp\left(\frac{-V(y_{t+1})}{\theta}\right)$$

= exp $\left[-.5\left(\epsilon_{t+1} - w_{t+1}\right)' \Sigma_t^{-1} \left(\epsilon_{t+1} - w_{t+1}\right) + .5\epsilon_{t+1}' \epsilon_{t+1} - .5\log \det \Sigma_t\right]$

and conditional entropy is given by 34

$$E\left(m_{t+1}\log m_{t+1}|y^t\right) = .5\left[|w_{t+1}|^2 + \operatorname{trace}\left(\Sigma_t - I\right) - \log \det \Sigma_t\right],$$

so that the worst case transition law is

$$y_{t+1} = Ay_t + Bu_t + C\hat{\epsilon}_{t+1} \tag{2.A.8}$$

where $\hat{\epsilon}_{t+1}$ is an i.i.d. Gaussian process with mean w_{t+1} and covariance matrix Σ_t .

Thus, it does not harm the malevolent agent to require that he chooses m_{t+1} so that the resulting worst-case model is Gaussian. Therefore, we can express the Bellman equation as

$$V(y_t) = \min_{w_{t+1}, \Sigma_t} \max_{u_t} \left\{ r(y_t, u_t) + \beta E V(y_{t+1}) + .5\beta \theta \left[|w_{t+1}|^2 + \operatorname{trace}\left(\Sigma_t - I\right) - \log \det \Sigma_t \right] \right\}$$
(2.A.9)

³³ To show this, let $\mu_t = Ay_t + Bu_t$ so that $y_{t+1} = \mu_t + C\epsilon_{t+1}$. Note that

$$\frac{-V\left(y_{t+1}\right)}{\theta} = \frac{\left(\epsilon_{t+1}'C'PC\epsilon_{t+1} + C'P\mu_t\right) + \left(\mu_t'P\mu_t + \rho\right)}{2\theta}.$$
(2.A.7)

Note that $-(\epsilon_{t+1} - w_{t+1})' \Sigma_t^{-1} (\epsilon_{t+1} - w_{t+1}) + \epsilon'_{t+1} \epsilon_{t+1} + \text{terms measurable w.r.t. } y^t = 0$ $-\epsilon'_{t+1}\Sigma_t^{-1}\epsilon_{t+1} - w'_{t+1}\Sigma_t^{-1}w_{t+1} + 2\epsilon'_{t+1}\Sigma_t^{-1}\epsilon_{t+1} + \epsilon_{t+1} = \epsilon'_{t+1}(I - \Sigma_t^{-1})\epsilon_{t+1} - 2\epsilon'_{t+1}\Sigma_t w_{t+1} - w'_{t+1}\Sigma_t^{-1}w_{t+1}$ Matching terms with the objects in (2.A.7) gives $I - \Sigma_t^{-1} = C'PC\theta$ and $w_{t+1} = (\theta I - C'PC)^{-1}C'P\mu_t$. 34 Conditional entropy $Em_{t+1} \log m_{t+1} | y_t$ evidently equals the expectation of

$$\left[-.5\left(\epsilon_{t+1}-w_{t+1}\right)'\Sigma_t^{-1}\left(\epsilon_{t+1}-w_{t+1}\right)+.5\epsilon_{t+1}'\epsilon_{t+1}-.5\log\det\Sigma_t\right]$$

under the distorted distribution formed by using m_{t+1} as a likelihood ratio with respect to the original $\mathcal{N}(0, I)$ distribution for ϵ_{t+1} . Thus, we have to evaluate $E_m[-.5 \operatorname{trace} \Sigma_t^{-1} (\epsilon_{t+1}$ $w_{t+1})(\epsilon_{t+1} - w_{t+1})' + .5\epsilon'_{t+1}\epsilon_{t+1}] - .5\log \det \Sigma_t. \text{ Using } \epsilon_{t+1} = w_{t+1} - (\epsilon_{t+1} - w_{t+1}),$ $\epsilon_{t+1}\epsilon'_{t+1} = w'_{t+1}w_{t+1} + (\epsilon_{t+1} - w_{t+1})'(\epsilon_{t+1} - w_{t+1}) + 2w'_{t+1}(\epsilon_{t+1} - w_{t+1}) \text{ and taking expectation}$ tations gives $E_m = w'_{t+1}w_{t+1} + \text{trace}\Sigma$. Therefore, $E_m[-.5\text{trace}\Sigma_t^{-1}(\epsilon_{t+1} - w_{t+1})(\epsilon_{t+1} - w_{t+1})(\epsilon$ $w_{t+1})' + .5\epsilon'_{t+1}\epsilon_{t+1}] - .5\log \det \Sigma_t = -.5 \operatorname{trace}(\Sigma - I) + .5w'_{t+1}w_{t+1} - .5\log \det \Sigma_t.$

54

subject to

$$y_{t+1} = Ay_t + Bu_t + C (w_{t+1} + \hat{\epsilon}_{t+1})$$
$$E\hat{\epsilon}_{t+1}\hat{\epsilon}'_{t+1} = \Sigma_t.$$

The minimizing choices of w_{t+1}, Σ_t are

$$w_{t+1} = Ky_t \tag{2.A.10}$$

$$\Sigma_t^{-1} = I - \theta^{-1} C' P C \tag{2.A.11}$$

and the maximizing choice of u_t is $u_t = -Fy_t$ where K satisfies (2.5.10), F satisfies (2.5.9), and where P is the pertinent fixed point of $T \circ \mathcal{D}$, where \mathcal{D} and T are defined in (2.5.6) and (2.5.8), and where θ in these formulas must be set to .5 times the θ in (2.A.8), because of the scaling we chose in the text of this chapter. Comparing these solutions with those in the text of this chapter tells us that the solution can be computed by solving the nonstochastic problem associated with the Bellman equation

$$V(y_t) = \min_{w_{t+1}} \max_{u_t} \left\{ r(y_t, u_t) + \beta E V(y_{t+1}) + .5\beta \theta |w_{t+1}|^2 \right\}.$$

With the same appropriate multiplication of θ by .5, formulas (2.5.9), (2.5.10), (2.5.6), (2.5.8) again give the pertinent P, K, and F. After P is computed, the minimizing Σ_t can be found from (2.A.11).

B. Matlab programs

A robust optimal linear regulator is defined by the system matrices Q, R, A, B, C, the discount factor β , and the risk-sensitivity parameter $\sigma \equiv -\theta^{-1}$. The Matlab program olrprobust.m implements the algorithm of section 2.7 by calling the optimal linear regulator program olrp.m. The program olrprobust solves a minimum problem, so that $\sigma < 0$ corresponds to a concern about robustness and R and Q should be approximately negative definite, where we say approximately because of the usual detectability qualifications. Call the program olrprobust as follows: [F,K,P,Pt]=olrprobust(beta,A,B,C,Q,R,sig). The objects returned by olrprobust determine the decision rule $u_t = -Fy_t$, the distortion $w_{t+1} = Ky_t$, the quadratic form in the value function -y'Py, and the distorted continuation value function $-y^{*'}(Pt)y^*$. The program doublex9 implements the doubling algorithm described in chapter 3 and by Hansen and Sargent (200XXXX(chapter 9)). To compute the robust rule with a discounted objective function, one has to induce doublex9 to solving a discounted problem by first setting $Ad = \sqrt{\beta}A, Bd = \sqrt{\beta}B$, calling [F,Kd,P,Pt]=doublex9(Ad,Bd,C,Q,R,sig), then finally setting $K = Kd/\sqrt{\beta}$. The program bayes4.m uses both olrprobust and doublex9 to compute robust decision rules and verifies that they give the same answers.

Part II Standard control and filtering

Chapter 3 Linear control theory

3.1. Introduction

This chapter analyzes the standard discounted linear-quadratic optimal control problem, called the optimal linear regulator. The robust decision maker to be described in later chapters adjusts this problem to reflect his doubts about the linear transition law. This chapter describes basic concepts of linear optimal control theory and efficient ways to compute solutions.¹ We describe methods that are faster than direct iterations on the Bellman equation (the Riccati equation) and are more reliable than solutions based on eigenvalue-eigenvector decompositions of the state-costate evolution equation.²

In later chapters, we use these techniques to formulate and solve various robust decision and estimation problems. Invariant subspace methods are key tools. In the present chapter, we show how they can be used to solve the Riccati equation that emerges from the Bellman equation for the linear regulator. In later chapters, we shall use invariant subspace methods in two important settings: (a) to compute robust decision rules and estimators in 'single agent' problems; and (b) to solve Ramsey problems in 'forward-looking' macroeconomic models. Invariant subspace methods also provide efficient algorithsms for analyzing and solving equilibria of rational expectations model that are formed by combining Euler equations and terminal conditions for a collection of decision makers with other equilibrium conditions and laws of motions for exogenous variables.

Section 3.2 decomposes the basic linear optimal control problem into subproblems that are more efficient to solve and describes classes of economic problems that give rise to such problems. Sections 3.3, 3.4, 3.5, and 3.6 describe

– 59 –

 $^{^1\,}$ Large parts of this chapter are based on Evan Anderson, Ellen McGrattan, and the authors (1996).

 $^{^2}$ Our survey of these methods draws heavily on Anderson (1978), Gardiner and Laub (1986), Golub, Nash and Van Loan (1979), Laub (1979,1991), and Pappas, Laub and Sandell (1980).

recent algorithms for solving these sub-problems. Subsection 3.4.2 briefly describes how to use invariant subspace methods to solve or approximately solve dynamic general equilibrium models.

3.2. Control problems

In this section, we pose three optimal control problems. We begin with a problem close to the much studied time-invariant deterministic optimal linear regulator problem. We label this the *deterministic regulator problem*. We then consider two progressively more general problems.

The first generalization introduces forcing sequences or "uncontrollable states" into the *deterministic regulator problem*. While this generalization is also a *deterministic regulator problem*, there are computational gains to exploiting the *a priori* knowledge that some components of the state vector are *uncontrollable*. We refer to this generalization as the *augmented regulator problem*. As we will see, a convenient first step for solving an *augmented regulator problem* is to solve a corresponding *deterministic regulator problem* in which the forcing sequence is "zeroed out." In other words, we obtain a piece of the solution to the *augmented regulator problem* by initially solving a problem with a smaller number of state variables.

The second generalization introduces, among other things, discounting and uncertainty into the *augmented regulator problem*. We refer to the resulting problem as the *discounted stochastic regulator problem*. Using well known transformations of the state and control vectors, we show how to convert this problem into a corresponding undiscounted *augmented regulator problem* without uncertainty. Therefore, while our original problem is a *discounted stochastic regulator problem*, we solve it by first solving a *deterministic regulator problem* with a smaller number of state variables, then solving a corresponding *augmented regulator problem*, and finally using this latter solution to construct the solution to the original problem in the manner described below.
3.2.1. Deterministic regulator problem

Choose a control sequence $\{v_t\}$ to maximize

$$-\sum_{t=0}^{\infty} \left(v_t' R v_t + y_t' Q_{yy} y_t \right),$$

subject to

$$y_{t+1} = A_{yy}y_t + B_y v_t$$
$$\sum_{t=0}^{\infty} \left(|v_t|^2 + |y_t|^2 \right) < \infty.$$
(3.2.1)

This control problem is a standard time-invariant, deterministic optimal linear regulator problem with one modification. We have added a stability condition, (3.2.1), that is absent in the usual formulation. This stability condition plays a central role in at least one important class of dynamic economic models: permanent income models. More will be said about these models later. In these models, the stability condition can be viewed as an infinite horizon counterpart to a terminal condition on the capital stock.

Following the literature on the time-invariant optimal linear regulator problem, we impose the following:

Definition: The pair (A_{yy}, B_y) is stabilizable if $y'B_y = 0$ and $y'A_{yy} = \lambda y'$ for some complex number λ and some complex vector y implies that $|\lambda| < 1$ or y = 0.

Assumption 1: (A_{yy}, B_y) is stabilizable.

Stabilizability is equivalent to the existence of a time-invariant control law that stabilizes the state (see Anderson and Moore, 1979, Appendix C). For our applications, it can often be verified by showing that a trivial control law, such as setting investment equal to zero, achieves this stability.

In solving this problem, we are primarily interested in specifications for which all of the state variables are "endogenous," and hence the following stronger restriction is met:

Definition: The pair (A_{yy}, B_y) is controllable if $y'B_y = 0$ and $y'A_{yy} = \lambda y'$ for some complex number λ and some complex vector y implies that y = 0. When (A_{yy}, B_y) is controllable, starting from an initialization of zero, the state vector can attain any arbitrary value in a finite number of time periods by an appropriate setting of the controls (see Anderson and Moore, 1979, Appendix C).³ For this reason, we can think of a state vector sequence with evolution equation governed by a pair (A_{yy}, B_y) that is controllable as being an *endogenous* state vector sequence.

While Assumption 1 gives us a nonempty constraint set, it is still possible that the *supremum* of the objective is not attained. We assume the following:

Assumption 2: The matrix Q_{yy} is positive semidefinite, and the matrix R is positive definite.

Among other things, this concavity assumption puts an upper bound of zero on the criterion function. Therefore, the *supremum* is finite (and nonpositive). We require that the *supremum* is attained.

Assumption 3: There exists a solution to the deterministic regulator problem for each initialization of y_0 .

A commonly used sufficient condition in the control theory literature for there to exist a solution is *detectability*. Factor $Q_{uu} = D_u D_u'$.

Definition: The pair (A_{yy}, D_y) is detectable if $D'_y y = 0$ and $A_{yy} y = \lambda y$ for some complex number λ and some complex vector y implies that $|\lambda| < 1$ or y = 0.

When the pair (A_{yy}, D_y) is detectable, it is optimal to choose a control sequence that stabilizes the state vector. In this case, the solution to the control problem is the same with or without the stability constraint (3.2.1). However, as we mentioned previously, for permanent income models the stability constraint is essential for obtaining an interpretable solution to the problem. For these models, detectability is too strong of a condition to impose. Chan, Goodwin and Sin (1984) give a weaker sufficient condition for there to exist a solution (see (iii) of Theorem 3.10). In the context of a continuous-time formulation, Hansen, Heaton and Sargent (1991) proposed a very similar sufficient condition for stabilizable systems based on a spectral representation of the *deterministic*

 $^{^3}$ This is one of five equivalent characterizations of *reachability* given in Appendix C of Anderson and Moore (1979). However, many other control theorists take one of these characterizations as the definition of *controllability*. For instance, see Kwakernaak and Sivan (1972) and Caines (1988). We choose to follow this latter convention.

regulator problem. Unfortunately, these conditions may be tedious to check in practice. Some of the solution algorithms we survey below could in principle be modified to detect a violation of Assumption 3.

A sufficient condition for convergence of one of the solution algorithms that we survey below is that the pair (A_{yy}, D_y) be *observable*:

Definition: The pair (A_{yy}, D_y) is observable if $D_y'y = 0$ and $A_{yy}y = \lambda y$ for some complex number λ and some complex vector y implies that y = 0.

Clearly, observability is stronger than detectability. Moreover, observability is guaranteed when the matrix Q_{yy} is nonsingular. When the pair (A_{yy}, D_y) is observable, the value function associated with the *deterministic regulator problem* is strictly concave in the state vector y (Caines and Mayne 1970, 1971).

The solution to the *deterministic regulator problem* takes the form

$$v_t = -F_y y_t$$

for some feedback matrix F_y . The stability constraint (3.2.1) guarantees that the eigenvalues of $A_{yy} - B_y F_y$ have absolute values that are strictly less than one because the state evolution equation when the optimal control is imposed is given by

$$y_{t+1} = (A_{yy} - B_y F_y) y_t.$$

3.2.2. Augmented regulator problem

Choose a control sequence $\{v_t\}$ to maximize

$$-\sum_{t=0}^{\infty} \left(v_t' R v_t + y_t' Q_{yy} y_t + 2 y_t' Q_{yz} z_t \right),$$

subject to

$$\begin{bmatrix} y_{t+1} \\ z_{t+1} \end{bmatrix} = \begin{bmatrix} A_{yy} & A_{yz} \\ 0 & A_{zz} \end{bmatrix} \begin{bmatrix} y_t \\ z_t \end{bmatrix} + \begin{bmatrix} B_y \\ 0 \end{bmatrix} v_t$$
$$\sum_{t=0}^{\infty} \left(|v_t|^2 + |y_t|^2 \right) < \infty.$$

We have modified the optimal linear regulator problem by including the *exogenous* forcing sequence $\{z_t\}$. The presumption here is that this partitioning may occur naturally in the specification of the original control problem. Of course, as is well known in the control theory literature, we could always transform an original state vector into controllable and uncontrollable components. Constructing this transformation, however, can be difficult to do in a numerically reliable way. In the next section we will display a class of optimal resource allocation problems associated with dynamic economies for which z_t contains a vector of taste and technology shifters. By assumption, this component of the state vector cannot be influenced by a control vector such as the level of investment.

For the *augmented regulator problem* to be well posed, we require that the forcing sequence be stable:

Assumption 4: The eigenvalues of A_{zz} have absolute values that are strictly less than one.

The solution to the *deterministic regulator problem* gives us a piece of the solution to the *augmented regulator problem*. More precisely, the solution to the augmented problem is

$$v_t = -F_y y_t - F_z z_t,$$

where the matrix F_y is the same as in the solution to the regulator problem for which the forcing sequence $\{z_t\}$ is zeroed out. Consequently, our solution methods entail first computing F_y by solving a deterministic regulator problem of lower dimension and then computing F_z given F_y .

3.2.3. Discounted stochastic regulator problem

Let $\{\mathcal{F}_t : t = 0, 1, ...\}$ denote an increasing sequence of sigma algebras (information sets) defined on an underlying probability space. We presume the existence of a "building block" process of conditionally homoskedastic martingale differences $\{\epsilon_t : t = 1, 2, ...\}$, which obeys

Assumption 5: The process $\{\epsilon_t : t = 1, 2, ...\}$ satisfies (i) $E(\epsilon_{t+1}|\mathcal{F}_t) = 0;$ (ii) $E(\epsilon_{t+1}\epsilon_{t+1}'|\mathcal{F}_t) = I.$ The discounted stochastic regulator problem is to choose a control process $\{u_t\}$, adapted to $\{\mathcal{F}_t\}$, to maximize

$$-E\left(\sum_{t=0}^{\infty}\beta^{t}\begin{bmatrix}u_{t}' & x_{t}'\end{bmatrix}\begin{bmatrix}R & W'\\W & Q\end{bmatrix}\begin{bmatrix}u_{t}\\x_{t}\end{bmatrix}\middle|\mathcal{F}_{0}\right),$$

subject to

$$x_{t+1} = Ax_t + Bu_t + C\epsilon_{t+1}$$
$$E\left(\sum_{t=0}^{\infty} \beta^t \left(|u_t|^2 + |x_t|^2\right) \middle| \mathcal{F}_0\right) < \infty.$$

The state vector x_t is taken to be the composite of the endogenous and exogenous state variables. Let $U_y = [I \ 0]$ be a matrix that *selects* the endogenous state vector $U_y x_t$ and $U_z = [0 \ I]$ be a matrix that *selects* the exogenous state vector $U_z x_t$ for an optimization problem with discounting. To justify our partitioning, the matrix A is restricted to satisfy $U_z A U_y' = 0$, and the matrix B is restricted to satisfy $U_z B = 0$. Notice that in addition to incorporating discounting and uncertainty, the *discounted stochastic regulator* includes cross-product terms between controls and states, captured with u'W'x, which are absent in the *augmented control problem*.

We now apply a standard trick for converting a discounted stochastic regulator problem to an augmented regulator problem. Using the well known certainty equivalence property of stochastic optimal linear regulator problems, we zero out the uncertainty without altering the optimal control law. That is, we are free to set the matrix C to zero and instead solve the resulting deterministic control problem. We eliminate discounting and cross-product terms between states and controls by using the transformations

$$y_t = \beta^{t/2} U_y x_t, \quad z_t = \beta^{t/2} U_z x_t, \quad v_t = \beta^{t/2} \left(u_t + R^{-1} W' x_t \right).$$

As is evident from these formulas, we have absorbed the discounting directly into the construction of the transformed state and control vectors. In addition, the cross-product matrix W is folded into the construction of the transformed control vector. We are left with a version of the *augmented regulator problem* with the following matrices:

$$\begin{bmatrix} A_{yy} & A_{yz} \\ 0 & A_{zz} \end{bmatrix} = \beta^{1/2} \left(A - BR^{-1}W' \right), \quad B_y = \beta^{1/2}U_y B,$$

$$\begin{bmatrix} Q_{yy} & Q_{yz} \\ Q_{yz'} & Q_{zz} \end{bmatrix} = Q - W R^{-1} W'.$$
(3.2.2)

Assumptions 1 - 4 are imposed on the constructed matrices on the left-hand side of the equal signs in (3.2.2).

As before, write the solution to the *augmented regulator problem* as

$$v_t = -F_y y_t - F_z z_t$$

Then the solution to the discounted stochastic regulator problem is

$$u_t = -Fx_t$$

where

$$F = \begin{bmatrix} F_y \\ F_z \end{bmatrix} + R^{-1}W'.$$

Also as before, the matrix F_y can be computed by solving the corresponding *deterministic regulator problem* with the forcing sequence "zeroed out." Subsequent sections will describe methods for computing F_y and F_z .

In macroeconomics, the discounted stochastic regulator problem is often obtained in the fashion of Kydland and Prescott (1982), who use it to replace a nonlinear-quadratic problem. Thus consider the nonquadratic optimization problem: choose an adapted (to $\{\mathcal{F}_t\}$) control process $\{u_t\}$ to maximize

$$-E\left(\sum_{t=0}^{\infty}\beta^{t}r\left(u_{t},x_{t}\right)\mid\mathcal{F}_{0}\right),$$
(3.2.3)

subject to

$$x_{t+1} = Ax_t + Bu_t + C\epsilon_{t+1}.$$

Here r is not required to be a quadratic function of u_t and x_t . When the associated constraints are nonlinear, sometimes we can substitute the nonlinear constraints into the criterion function to obtain a problem of the form of (3.2.3). Kydland and Prescott (1982) simply replace the function r by a quadratic form in $[u_t' x_t']'$ as required for the *discounted stochastic regulator problem*, where the quadratic function is designed to "approximate" r well near a particular value for the state vector.⁴ In chapter 5, we describe a different approach

⁴ While Kydland and Prescott (1982) apply an *ad hoc* global approximation to r in which the range of approximation is adapted to the amount of underlying uncertainty, many later researchers have instead simply used a local Taylor series approximation around some "nonstochastic" steady state produced by shutting down all randomness in the model. Kydland and Prescott (1982) note that for the range of uncertainty they considered, the two methods gave similar answers. In forming the linear quadratic problem, it is important to substitute the non-linear constraints into the objective function before taking a Taylor series approximation.

where, by design, the initial optimal resource allocation problem can be directly converted into a *discounted stochastic regulator problem*.

3.3. Solving the deterministic linear regulator problem

In this section we describe ways to solve for the matrix F_y . Recall that this matrix has a double role. First, it gives the control law for a particular deterministic regulator problem. More importantly for us, it also gives a piece of the solution to the discounted stochastic regulator problem.

In describing methods for computing F_y , it is convenient to work with the state-costate equations associated with the Lagrangian

$$\mathcal{L} = -\sum_{t=0}^{\infty} \left[y_t' Q_{yy} y_t + v_t' R v_t + 2\mu_{t+1}' \left(A_{yy} y_t + B_y v_t - y_{t+1} \right) \right].$$
(3.3.1)

First-order necessary conditions for the maximization of \mathcal{L} with respect to $\{v_t\}_{t=0}^{\infty}$ and $\{y_t\}_{t=0}^{\infty}$ are

$$v_t: \quad Rv_t + B_y' \mu_{t+1} = 0, \quad t \ge 0$$
 (3.3.2)

$$y_t: \qquad \mu_t = Q_{yy}y_t + A_{yy}'\mu_{t+1}, \qquad t \ge 0.$$
 (3.3.3)

To obtain a composite state-costate evolution equation, solve (3.3.2) for v_t , substitute the solution into the state evolution equation, and stack the resulting equation and (3.3.3) and write the state-costate evolution equation as

$$L\begin{bmatrix} y_{t+1}\\ \mu_{t+1} \end{bmatrix} = N\begin{bmatrix} y_t\\ \mu_t \end{bmatrix}, \qquad (3.3.4)$$

where

$$L \equiv \begin{bmatrix} I & B_y R^{-1} B_y' \\ 0 & A_{yy'} \end{bmatrix}, N \equiv \begin{bmatrix} A_{yy} & 0 \\ -Q_{yy} & I \end{bmatrix}.$$

For a continuous-time system the a corresponding differential equation for states and costates is

$$\begin{bmatrix} Dy_t \\ D\mu_t \end{bmatrix} = H \begin{bmatrix} y_t \\ \mu_t \end{bmatrix}, \qquad (3.3.5)$$

where

$$H \equiv \begin{bmatrix} A_{yy} & -B_y R^{-1} B'_y \\ -Q_{yy} & -A_{yy}' \end{bmatrix}, \qquad (3.3.6)$$

which assembles the first-order conditions for the problem with criterion $-\int_0^\infty [y(t)'Q_{yy}y(t)+u(t)'Ru(t)]dt$ and law of motion $Dy(t) = A_{yy}y(t) + B_yu(t)$, where D is the time-differentiation operator. We describe several methods for solving equations (3.3.4) and (3.3.5). Formally, we will devote most of our attention to the discrete-time system (3.3.4). As we will see, methods designed for solving the continuous-time system (3.3.5) can be adapted easily to solve the discrete-time system (3.3.4), and conversely.

We want the solution of (3.3.4) that stabilizes the state-costate vector sequence for any initialization y_0 . Since we have transformed the state vector to eliminate discounting, we impose stability in the form of square summability:

$$\sum_{t=0}^{\infty} \left| \begin{bmatrix} y_t \\ \mu_t \end{bmatrix} \right|^2 < \infty, \tag{3.3.7}$$

for the discrete-time system (3.3.4). (We impose the analogous square integrability restriction on the continuous time system (3.3.5)).

One way to ascertain the solution to the *deterministic regulator problem* is to find an *initial* costate vector expressed as a function of the initial state vector y_0 that guarantees the stability of system (3.3.4) or (3.3.5). The initialization of the costate vector takes the form $\mu_0 = P_y y_0$ and is replicated over time. Substituting $P_y y_t$ for μ_t into (3.3.4), we find that

$$(I + B_y R^{-1} B_y' P_y) y_{t+1} = A_{yy} y_t A_{yy}' P_y y_{t+1} = -Q_{yy} y_t + P_y y_t.$$
(3.3.8)

It is straightforward to verify that

$$\left(I + B_y R^{-1} B_y' P_y\right)^{-1} = I - B_y \left(R + B_y' P_y B_y\right)^{-1} B_y' P_y.$$
(3.3.9)

Solving the first equation in (3.3.8) for y_{t+1}

$$y_{t+1} = (A_{yy} - B_y F_y) y_t, \qquad (3.3.10)$$

where

$$F_{y} \equiv \left(R + B_{y}' P_{y} B_{y}\right)^{-1} B_{y}' P_{y} A_{yy}.$$
 (3.3.11)

Premultiplying (3.3.10) by $A'_{yy}P_y$ gives

$$A'_{yy}P_{y}y_{t+1} = \left(A'_{yy}P_{y}A_{yy} - A'_{yy}P_{y}B_{y}F_{y}\right)y_{t}.$$
(3.3.12)

For the right-hand side of equation (3.3.12) to agree with the right-hand side of the second equation of (3.3.8) for any initialization y_0 , it must be that

$$P_{y} = Q_{yy} + A'_{yy}P_{y}A_{yy} - A'_{yy}P_{y}B_{y}\left(R + B'_{y}P_{y}B_{y}\right)^{-1}B'_{y}P_{y}A_{yy},$$

= $Q_{yy} + (A_{yy} - B_{y}F_{y})'P_{y}\left(A_{yy} - B_{y}F_{y}\right) + F'_{y}RF_{y}$, (3.3.13)

which is the familiar *Riccati equation*. In other words, the matrix P_y used to set the initial condition on the costate vector is also a solution to the Riccati equation (3.3.13). With this initialization, the costate relation $\mu_t = P_y y_t$ holds for all $t \ge 0$. Finally, it follows from (3.3.10) that this state-costate solution is implemented by the control law $v_t = -F_y y_t$.

The remainder of this section is organized as follows. In the first subsection, we initially consider the case in which the matrix A_{yy} is nonsingular. While this case is studied for pedagogical simplicity, it is also of interest in its own right. In the second subsection, we then treat the more general case in which A_{yy} can be singular. As emphasized by Pappas, Laub and Sandell (1980), singularity in A_{yy} occurs naturally in dynamic systems with delays. One of our example economies used in our numerical experiments has a singular matrix A_{yy} . Finally, in the third subsection we study the continuous-time counterpart to the *deterministic regulator problem*. We describe an alternative solution method and show how to convert a discrete-time regulator problem into a continuous-time regulator with the same relation between optimally chosen state and co-state vectors. We defer the discussion of the numerical algorithms used for implementing these methods until the next section.

3.3.1. Nonsingular A_{yy}

When the matrix A_{yy} is nonsingular, we can solve (3.3.4) for $\begin{bmatrix} y_{t+1} \\ \mu_{t+1} \end{bmatrix}$:

$$\begin{bmatrix} y_{t+1} \\ \mu_{t+1} \end{bmatrix} = M \begin{bmatrix} y_t \\ \mu_t \end{bmatrix}, \qquad (3.3.14)$$

where

$$M \equiv L^{-1}N = \begin{bmatrix} A_{yy} + B_y R^{-1} B'_y A'_{yy}^{-1} Q_{yy} & -B_y R^{-1} B'_y A'_{yy}^{-1} \\ -A'_{yy}^{-1} Q_{yy} & A'_{yy}^{-1} \end{bmatrix}.$$
 (3.3.15)

We find the matrix P_y by locating the stable *invariant subspace* of the matrix M.

Definition: An invariant subspace of a matrix M is a linear space C of possibly complex vectors for which MC = C.

Invariant subspaces are constructed by taking linear combinations of eigenvectors of M. A stable invariant subspace is one for which the corresponding eigenvalues have absolute values less than one. To solve the model, we seek a matrix P_y such that $\begin{bmatrix} I \\ P_y \end{bmatrix} y$ is in the stable invariant subspace of M for every n dimensional vector y. We now elaborate on how to compute this subspace.

The matrix M has a particular structure that we can exploit in characterizing its eigenvalues. To represent this structure, we introduce a matrix J given by

$$J \equiv \begin{bmatrix} 0 & -I \\ I & 0 \end{bmatrix}.$$

Notice that $J^{-1} = J' = -J$.

Definition: A matrix M is symplectic if MJM' = J.

It is straightforward to verify that M given by (3.3.15) is symplectic. It follows that

$$M' = J^{-1}M^{-1}J. (3.3.16)$$

Therefore, the transpose of M is similar to its inverse. Recall that similar matrices define the same linear transformation but with respect to a different coordinate system. Thus M' and M^{-1} share the same eigenvalues. For any matrix M, the eigenvalues of M^{-1} are the reciprocals of the eigenvalues of M, so it follows that the eigenvalues of a real symplectic matrix come in reciprocal pairs, and the number of stable eigenvalues cannot exceed the number of states n. However, merely requiring M to be symplectic permits there to be eigenvalues with absolute values equal to one, and so we will need an additional argument to show that there are exactly n stable eigenvalues.

To locate the stable invariant subspace of the symplectic matrix M, we follow Laub (1979) and (block) triangularize M:

$$V^{-1}MV = W$$
$$W = \begin{bmatrix} W_{11} & W_{12} \\ 0 & W_{22} \end{bmatrix}, \qquad (3.3.17)$$

where V is a nonsingular matrix. By construction, the matrices M and W are similar. The matrix partitions in (3.3.17) are built to coincide with the number of stable and unstable eigenvalues. In particular, the absolute values of the eigenvalues of W_{11} are stable.

A special case of this decomposition is an appropriately ordered Jordan decomposition of M as was used by Vaughan (1970) in developing an invariant subspace algorithm for computing P_y . Laub (1991) traces this solution strategy back to the 19th century and credits MacFarlane (1963) and Potter (1966) with introducing it to the control literature. As emphasized by Laub (1991), it is preferable to build algorithms based on other upper triangular decompositions that are more stable numerically. The Jordan decomposition is particularly problematic when the symplectic matrix M has eigenvalues with multiplicities greater than one (see also Golub and Wilkinson 1976). In the next section, we describe alternative Schur decompositions, which are more reliable numerically.

To use this triangularization to calculate P_y , apply V^{-1} to both sides of the state equation (3.3.14):

$$\check{y}_{t+1} = W\check{y}_t,$$

where

$$\check{y}_t = V^{-1} \begin{bmatrix} y_t \\ \mu_t \end{bmatrix}$$

This transformation permits us to study asymptotic properties in terms of two smaller uncoupled subsystems. Partition \check{y}_t into two blocks with dimensions given by the number of stable and unstable eigenvalues:

$$\check{y}_t \equiv \begin{bmatrix} \check{y}_{1,t} \\ \check{y}_{2,t} \end{bmatrix}.$$

Then

$$\check{y}_{2,t+1} = W_{22}\check{y}_{2,t+1}$$

and the solution sequence $\{\check{y}_{2,t}\}$ fails to converge to zero unless it is initialized at zero. Setting $\check{y}_{2,0}$ at zero can be accomplished by an appropriate initialization of the costate vector, as we now verify.

Partition the matrices V and V^{-1} as

$$V = \begin{bmatrix} V_{11} & V_{12} \\ V_{21} & V_{22} \end{bmatrix}, V^{-1} = \begin{bmatrix} V^{11} & V^{12} \\ V^{21} & V^{22} \end{bmatrix}.$$

Since V is nonsingular and there exists a (stable) solution to the optimal control problem, we must have

$$V^{21}y_t + V^{22}\mu_t = 0. ag{3.3.18}$$

The rank of the matrix $\begin{bmatrix} V^{21} & V^{22} \end{bmatrix}$ equals the number of unstable eigenvalues of M, and thus the rank of its null space must equal the number of stable eigenvalues. For a solution to exist for every initialization $y_0 = y$, it follows from (3.3.18) that there must exist a μ such that

$$V^{21}y + V^{22}\mu = 0.$$

Thus the dimensionality of the null space of $\begin{bmatrix} V^{21} & V^{22} \end{bmatrix}$ must also be at least n. Therefore, M has exactly n stable eigenvalues, and the matrix partition V^{22} is nonsingular. Solving (3.3.18) for μ_t gives

$$\mu_t = -\left(V^{22}\right)^{-1} V^{21} y_t.$$

Consequently, the matrix P_y used to initialize the costate vector is given by

$$P_y = -(V^{22})^{-1} V^{21} = V_{21} V_{11}^{-1}, \qquad (3.3.19)$$

where the second equality follows from the fact that the rank of $\begin{bmatrix} V_{11} \\ V_{21} \end{bmatrix}$ is n, and

$$\begin{bmatrix} V^{21} & V^{22} \end{bmatrix} \begin{bmatrix} V_{11} \\ V_{21} \end{bmatrix} = 0.$$

3.3.2. Singular A_{yy}

We now extend the solution method to accommodate singularity in A_{yy} . This method avoids inverting the L matrix in (3.3.4). Instead of locating the stable invariant subspace of M, a deflating subspace method finds the stable deflating subspace of the pencil $\lambda L - N$.

Definition: A pencil $\lambda L - N$ is the family of matrices $\{\lambda L - N\}$ indexed by the complex variable λ .

Definition: A deflating subspace of the pencil $\lambda L - N$ is a subspace C of complex vectors such that the dimension of C is at least as large as the dimension of the sum of the subspaces LC and NC.

For the matrices L and N of equation (3.3.4), it can be verified that the intersection of their null spaces contains only the zero vector.⁵ This ensures us that a generalized eigenvalue problem is well posed. When a subspace C is deflating, there exists a vector y in C that solves the generalized eigenvalue problem

$$\lambda Ly = Ny$$

(see Theorem 2.1 in Stewart 1972). Implicitly, we are including the possibility of a solution with $\lambda = \infty$, which occurs when y is in the null space of L but not in the null space of N. As with the previous (invariant subspace) method, the deflating subspace of interest for solving the optimal control problem is the deflating subspace associated with the stable state-costate sequence. The stable deflating subspace is the subspace associated with the stable generalized eigenvectors (the eigenvectors associated with generalized eigenvalues with absolute values strictly less than one.) Hence we solve the model by finding a matrix P_y such that $\begin{bmatrix} I \\ P_y \end{bmatrix} y$ is in the stable deflating subspace of the pencil $\lambda L - N$.

Recall that when A_{yy} is nonsingular, the matrix M is symplectic. More generally, system (3.3.4) is associated with a symplectic pencil

Definition: A pencil $\lambda L - N$ is symplectic if LJL' = NJN'.

Pappas, Laub and Sandell (1980, Theorem 4) show that the generalized eigenvalues of the symplectic pencil $(\lambda L - N)$ come in reciprocal pairs, just as the eigenvalues of M do when A_{yy} is nonsingular. Hence we again have that the number of stable generalized eigenvalues is no greater than n. Furthermore, we can imitate our argument in the case in which A_{yy} is nonsingular to show that there are exactly n stable generalized eigenvalues.⁶

We triangularize the state-costate system (3.3.4) using the solutions to the generalized eigenvalue problem. As in Theorem 2.1 of Stewart (1972), there

⁵ See Theorem 3 of Pappas, Laub and Sandell (1980) for the case in which (A_{yy}, D_y) is detectable. As we noted previously, the restriction to a detectable system rules out some interesting economic models. More generally, nonexistence of a common nonzero vector in the null spaces of N and L can be shown by way of contradiction. Suppose there is a common nonzero vector in the null space. Then the matrix $(I + Q_{yy}B_yR^{-1}B'_y)$ is singular. However, this singularity contradicts Theorem 1 of Kimura (1988).

⁶ Theorems 3 and 4 of Pappas, Laub and Sandell (1980) establish this result when the pair (A_{yy}, D_y) is detectable.

exists a decomposition of the pencil $\lambda L - N$ such that

$$ULV = T = \begin{bmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{bmatrix}, UNV = W = \begin{bmatrix} W_{11} & W_{12} \\ 0 & W_{22} \end{bmatrix},$$
(3.3.20)

where U and V are unitary matrices and the matrix partitions have the same number, n, of elements as the number of entries in the state vector y_t . Premultiplication of the pencil $\lambda L - N$ by the nonsingular matrix U preserves the solutions to the generalized eigenvalue problem, and postmultiplication by V alters the generalized eigenvectors but not the eigenvalues. A consequence of the triangularization is that the solutions to the generalized eigenvalue problem for the original system are constructed directly from the solutions to the following two smaller problems: $\lambda T_{11}\tilde{u} = W_{11}\tilde{u}$

$$\lambda T_{11}y = W_{11}y$$

 $\lambda T_{22}\tilde{y} = W_{22}\tilde{y}.$
(3.3.21)

As with the invariant subspace method, we build the blocks of the triangularization so that the generalized eigenvalues of the first problem in (3.3.21) satisfy $|\lambda| < 1$, and for the second problem $|\lambda| > 1$. As a consequence, the span of the first *n* columns of *V* gives the vectors of the deflating subspace we seek. The span of the remaining *n* columns contains the problematic initializations of the state-costate vector for which the implied sequence of state-costate vectors diverges exponentially. In addition, it includes the span of the generalized eigenvectors associated with infinite eigenvalues. Imitating the solution method when A_{yy} is nonsingular, we initialize the costate vector as $\mu_t = P_y y_t$, where the matrix P_y is again given by (3.3.19).

To understand better the nature of this unstable subspace, recall that an eigenvector associated with an infinite eigenvalue is in the null space of T_{22} . Suppose the triangularization of L and N is built so that we can further partition the matrices:

$$T_{22} = \begin{bmatrix} M_{11} & M_{12} \\ 0 & 0 \end{bmatrix}$$
$$W_{22} = \begin{bmatrix} O_{11} & O_{12} \\ 0 & O_{22} \end{bmatrix},$$

where the matrices M_{11} and O_{22} are nonsingular. Such a triangularization always exists. Consider solving the following equation recursively for a sequence $\{\tilde{y}_{t+1}\}$; for each t solve for \tilde{y}_{t+1} given \tilde{y}_t by using

$$T_{22}\tilde{y}_{t+1} = W_{22}\tilde{y}_t.$$

For this equation to have a solution, the second component of \tilde{y}_t must be zero for all t because

$$O_{22}\tilde{y}_{t,2} = 0, \tag{3.3.22}$$

and O_{22} is nonsingular. In addition to eliminating the nonexistence problem, imposing this restriction also resolves the multiplicity problem. Note that the multiplicity problem for the triangular system is that for a given t, (3.3.22) does not restrict $\tilde{y}_{t+1,2}$. However, applying (3.3.22) at t+1 resolves the problem.

3.3.3. Continuous-time systems

To conclude this section, we consider solving continuous-time Hamiltonian systems of the form (3.3.5). The defining feature of a *Hamiltonian* matrix is:

Definition: A matrix H is Hamiltonian if JH is symmetric.

The matrix H in (3.3.5), (3.3.6) clearly satisfies this property. It follows that

$$H' = -JHJ^{-1}$$

which in turn implies that the matrix H' is similar to -H. Consequently, the eigenvalues of a real Hamiltonian matrix come in pairs that are symmetric about the imaginary axis of the complex plane. The *stable* eigenvalues of a Hamiltonian matrix are those whose real parts are strictly negative. Similar arguments to those given above guarantee that there are exactly n stable eigenvalues of H. Therefore, (3.3.5) can be solved by using an invariant subspace method and its associated decomposition (3.3.17), provided that the classification of stable and unstable eigenvalues is modified appropriately.⁷

There is an alternative approach for solving a continuous-time Hamiltonian system. Given a Hamiltonian matrix H, another Hamiltonian matrix G is constructed with the same stable and unstable invariant subspaces. The matrix

⁷ Deflating subspace methods are not needed for solving the class of continuous-time quadratic control problems considered here because we can form directly the Hamiltonian matrix and apply an invariant subspace method. However, as we have formulated it, the continuous-time problem does not permit systems with finite gestation lags in making investment goods productive or systems for which consumption services depend on only a finite interval of past consumptions.

G is called the "sign" of the matrix H, and is defined as follows. Take the Jordan decomposition of H:

$$H = V \begin{bmatrix} \Lambda_{11} & 0\\ 0 & \Lambda_{22} \end{bmatrix} V^{-1},$$

where Λ_{11} is an upper triangular matrix with the eigenvalues of H that have strictly negative real parts on the diagonals, and Λ_{22} is an upper triangular matrix with the eigenvalues of H that have strictly positive real parts on the diagonals. Then

$$G = \operatorname{sign}(H) \equiv V \begin{bmatrix} -I & 0 \\ 0 & I \end{bmatrix} V^{-1}.$$

Thus the sign of a matrix is a new matrix with the same eigenvectors as the original matrix and with eigenvalues replaced by -1 or 1 depending on the signs of the real parts of the original eigenvalues.

The matrix P_y can be inferred directly from G. To see this, we use an insight from Roberts (1980). By construction, all of the stable eigenvalues of G are equal to -1. Consequently, the matrix P_y satisfies the following eigenvalue problem:

$$G\begin{bmatrix}I\\P_y\end{bmatrix}y = -\begin{bmatrix}I\\P_y\end{bmatrix}y$$

for any n dimensional vector y, and the matrix P_y solves the affine equation

$$G\begin{bmatrix}I\\P_y\end{bmatrix} + \begin{bmatrix}I\\P_y\end{bmatrix} = 0.$$
(3.3.23)

This method is implemented by finding fast ways to compute the "sign" of a matrix.

While the matrix sign method is directly applicable for solving continuoustime Hamiltonian systems, Hitz and Anderson (1972) and Gardiner and Laub (1986) show how to use it to locate deflating subspaces of discrete-time systems. Consider the generalized eigenvalue problem for the symplectic pencil

$$\lambda Ly = N$$

Then

$$(1+\lambda) (L-N) y = (1-\lambda) (L+N) y.$$

Since the only common vector in the null space of L and N is zero, we construct the solution to the eigenvalue problem

$$\delta y = (L - N)^{-1} (L + N) y,$$

where

$$\delta = \frac{1+\lambda}{1-\lambda}.$$

Consequently, the stability relations (3.2.1) carry over here as well, and we apply the matrix sign algorithm to $(L - N)^{-1}(L + N)$.

It also turns out that $(L - N)^{-1}(L + N)$ is a Hamiltonian matrix, which we can exploit in computation. To verify the Hamiltonian structure, note that

$$(L - N) J (L' + N') = LJL' - NJN' - NJL' + LJN'$$

= -NJL' + LJN'
= NJN' - LJL' - NJL' + LJN''
= - (L + N) J (L' - N')

where we have used the fact that $\lambda L - N$ is a symplectic pencil. Therefore,

$$J(L-N)^{-1}(L+N) = (L'+N')(L'+N')^{-1}J(L-N)^{-1}(L+N)$$

= $(L'+N')[-(L-N)J(L'+N')]^{-1}(L+N)$
= $(L'+N')[(L+N)J(L'-N')]^{-1}(L+N)$
= $(L'+N')(L'-N')^{-1}J',$

which proves that $(L - N)^{-1}(L + N)$ is a Hamiltonian matrix.

In summary, by construction, the stable (unstable) invariant subspace of the Hamiltonian matrix $(L-N)^{-1}(L+N)$ coincides with the stable (unstable) deflating subspace of the symplectic pencil $\lambda L - N$. This coincidence permits us to compute the matrix P_y used for initializing the costate vector for the discretetime system (3.3.4) by applying a matrix sign algorithm to $(L-N)^{-1}(L+N)$.

3.4. Computational techniques for solving Riccati equations

We consider three types of algorithms for computing P_y :

- (i) Schur algorithm;
- (ii) doubling algorithm;
- (iii) matrix sign algorithm.

A Schur algorithm is based on locating a stable subspace using a *Schur decomposition* of the state-costate system. As we noted in the previous section, once a stable subspace is located, the relevant Riccati equation solution P_y is easily computed. There are two versions of a Schur decomposition, depending on whether the matrix A_{yy} is known to be nonsingular or not. A Schur decomposition gives a more reliable way of locating stable spaces than the familiar Jordan decomposition and its generalization for pencils.

A doubling algorithm is an iterative method for speeding up the dynamic programming Riccati equation iteration by *doubling* the number of time periods in each iteration.

Recall from our discussion in the previous section that the stable deflating subspace of the pencil $\{\lambda L - N\}$ coincides with the invariant subspace of the sign of the matrix $(L - N)^{-1}(L + N)$ associated with the eigenvalue -1. A matrix sign algorithm is an iterative method for computing the sign of $(L-N)^{-1}(L+N)$ from which we can recover P_y easily. See section 3.4.4 for details of the matrix sign algorithm.

3.4.1. Schur algorithm

Suppose the matrix A_{yy} is nonsingular. As we noted in section 3, the matrix P_y can be found by locating the stable invariant subspace of the matrix M given in (3.3.15). In some of our numerical calculations, we use what is referred to as a *real* Schur decomposition of M to locate its invariant subspace.

Definition: The real Schur decomposition of a real matrix M is an orthogonal matrix \hat{V} and a real upper block triangular matrix \hat{W} such that

$$\hat{V}'M\hat{V} = \hat{W} = \begin{bmatrix} \hat{W}_{11} & \hat{W}_{12} & \dots & \hat{W}_{1m} \\ 0 & \hat{W}_{22} & \dots & \hat{W}_{2m} \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & \hat{W}_{mm} \end{bmatrix}$$

where \hat{W}_{ii} is either a scalar or a 2 × 2 matrix with complex conjugate eigenvalues.⁸

A real Schur decomposition is a computationally convenient version of the block triangular decomposition (3.3.17) used to compute P_y when A_{yy} is nonsingular. Golub and Van Loan (1989) describe how to compute the real Schur decomposition (in particular, see sections 7.4 and 7.5). Recall that the block triangular matrix W in (3.3.17) results from partitioning the eigenvalues into stable and unstable eigenvalues. Algorithms that compute the real Schur decomposition of a matrix typically do not partition the diagonal blocks of \hat{W} according to stability. Instead, given an arbitrary real Schur decomposition $M = \hat{V}\hat{W}\hat{V}'$, one can use the approaches described in either Bai and Demmel (1993) or Stewart (1976) to construct a sequence of orthogonal transformations that reorder the diagonal blocks of \hat{W} , while updating \hat{V} so that $M = \hat{V}\hat{W}\hat{V}'$ holds at every step.

In summary, the steps for implementing a Schur algorithm are

- (1) form the matrix M in (3.3.15);
- (2) form a real Schur decomposition of M where the first n columns of \hat{V} , written in a partitioned form as $[\hat{V_{11}}' \quad \hat{V_{21}}']'$, are a basis for the stable invariant subspace of M;
- (3) solve $P_y \hat{V}_{11} = \hat{V}_{21}$ for P_y .

We recommend computing the real Schur decomposition of M by using the LAPACK function DGEES; P_y in step (3) can be computed using the built-in MATLAB operator '/', which solves a linear equation using Gaussian elimination with partial pivoting.

⁸ There is also a *complex Schur decomposition* of a real or complex matrix in which \hat{V} is a unitary matrix and \hat{W} is upper triangular.

A deflating subspace method is required when A_{yy} is singular and likely to be more stable numerically when A_{yy} is nearly singular. To implement this approach in practice, we use an ordered real generalized Schur decomposition to find an appropriate triangularization of the state-costate dynamical system (see Van Dooren (1982)).

Definition: A generalized real Schur decomposition of a real matrix pencil $\lambda L - N$ is a pair of orthogonal matrices \hat{U} and \hat{V} , a real upper triangular matrix \hat{T} , and a real upper block triangular matrix \hat{W} , such that

$$\hat{U}L\hat{V} = \hat{T} = \begin{bmatrix} \hat{T}_{11} & \hat{T}_{12} & \dots & \hat{T}_{1m} \\ 0 & \hat{T}_{22} & \dots & \hat{T}_{2m} \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & \hat{T}_{mm} \end{bmatrix}$$
$$\hat{U}N\hat{V} = \hat{W} = \begin{bmatrix} \hat{W}_{11} & \hat{W}_{12} & \dots & \hat{W}_{1m} \\ 0 & \hat{W}_{22} & \dots & \hat{W}_{2m} \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & \hat{W}_{mm} \end{bmatrix}$$

where the pencil $\lambda \hat{T}_{ii} - \hat{W}_{ii}$ is either a 1 × 1 matrix pencil or a 2 × 2 matrix pencil with complex conjugate generalized eigenvalues.

As with the real Schur decomposition, we initially compute a generalized real Schur decomposition of $\lambda L - N$ without regard to whether the generalized eigenvalues are stable or not. We then reorder the diagonal blocks of \hat{T} and \hat{W} so that the generalized eigenvalues are partitioned in the manner required by (3.3.20). This partitioning can be done using the algorithms described in Van Dooren (1981,1982) or in Kågström and Poromaa (1994).

Thus the steps for implementing a generalized Schur algorithm are

- (1) form the matrices L and N in (3.3.4);
- (2) form a generalized real Schur decomposition of the pencil $\lambda L N$ where the first *n* columns of \hat{V} , written in a partitioned form as $[\hat{V_{11}}' \quad \hat{V_{21}}']'$, span the deflating subspace of the pencil $\lambda L N$;
- (3) solve $P_y \hat{V}_{11} = \hat{V}_{21}$ for P_y .

3.4.2. Digression: solving DGE models with distortions

Linear or log-linear approximations to the equilibrium conditions of dynamic general equilibrium (DGE) models take one of the forms

$$Ly_{t+1} = Ny_t + \tilde{G}z_t \tag{3.4.1}$$

or, if L is nonsingular,

$$y_{t+1} = My_t + Gz_t (3.4.2)$$

where $M = L^{-1}N$ and z_t is a vector of forcing variables governed by a law of motion

$$z_{t+1} = A_{22} z_t, \tag{3.4.3}$$

where the eigenvalues of A_{22} are all less than or equal to unity in modulus.⁹ We shall consider the case in which L is nonsingular. We assume that the eigenvalues of M split into equal numbers of stable and unstable ones so that we can obtain a real Schur decomposition of $M = V^{-1}MV = W = \begin{bmatrix} W_{11} & W_{12} \\ 0 & W_{22} \end{bmatrix}$ where W_{11} is a stable matrix and W_{22} is an unstable matrix. The assumption that the eigenvalues split in this way is tantamount to assuming that there exists a unique stabilizing solution of (3.4.1).

Using $M = VWV^{-1}$ in (3.4.2) and premultiplying both sides by V^{-1} gives

$$V^{-1}y_{t+1} = WV^{-1}y_t + V^{-1}Gz_t (3.4.4)$$

or

$$y_{t+1}^* = Wy_t^* + G^* z_t \tag{3.4.5}$$

where $y_t^* = V^{-1}y_t$ and $G^* = V^{-1}G$. Express (3.4.5) in terms of the uncoupled dynamic system

$$y_{1t+1}^* = W_{11}y_{1t}^* + W_{12}y_{2t}^* + G_1^* z_t aga{3.4.6a}$$

$$y_{2t+1}^* = W_{22}y_{2t}^* + G_2^*z_t. aga{3.4.6b}$$

⁹ This assumption can be relaxed to be that the eigenvalue of maximum modulus of A_{22} times the reciprocal of the largest eigenvalue of A_{22} is strictly less than unity. Tom: check the sign of this statement.

Where \tilde{L} is the lag operator, rewrite (3.4.6b) as $(I - W_{22}\tilde{L})y_{2t+1}^* = G_2^* z_t$ or $-W_{22}\tilde{L}(I - W_{22}^{-1}\tilde{L}^{-1})y_{2t+1}^* = G_2^* z_t$ or ¹⁰

$$y_{2t}^* = -W_{22}^{-1} \left(I - W_{22}^{-1} \tilde{L}^{-1} \right)^{-1} G^* z_t.$$
(3.4.7)

Substituting this into (3.4.6a) and rearranging gives

$$y_{t+1}^* = W_{11}y_{1t}^* + \left[G_1^* - W_{12}W_{22}^{-1}\left(I - W_{22}^{-1}\tilde{L}^{-1}\right)^{-1}G_2^*\right]z_t.$$
 (3.4.8)

Equations (3.4.7), (3.4.8) give the stabilizing solution for the uncoupled dynamic system cast in terms of y_t^* . To retrieve the original variables, we simply use $y_t = Vy_t^*$.

The very same solution would also be sustained as the solution of the stochastic system in which (3.4.3) is replaced by the stochastic law of motion

$$z_{t+1} = A_{22}z_t + Cw_{t+1} \tag{3.4.9}$$

where w_{t+1} is a martingale difference sequence with identity covariance matrix; and where y_{t+1} on the left side of (3.4.1) and (3.4.2) is replaced by $E[y_{t+1}|y_t, z^t]$, where here E is the mathematical expectation operator and z^t denotes the history of the z_s process up to and including t. Equations (3.4.7), (3.4.8) are also the heart of the solution that would obtain if were we to assume that in a stochastic system the state z_t is not observed, but that noisy signals Y_t that are linearly related to it. In that case, the solution is to replace z_t in (3.4.7), (3.4.8) with $E[z_t|Y^t]$. The projection $E[z_t|Y^t]$ can be computed recursively using the standard Kalman filtering formulas reported in chapter 4.

¹⁰ These formulas can be viewed as extensions to the vector case of formulas found in Sargent (1987a, ch IX, pp. ???).

3.4.3. Doubling algorithm

Dynamic programming solves the infinite horizon problem by backward induction, which leads to iterations on the Riccati equation (3.3.13). A doubling algorithm accelerates this approach. It preserves the idea of approximating the solution to the infinite horizon problem by a sequence of finite horizon problems, but instead of increasing the horizon by one time period in each iteration, the number of time periods gets *doubled*.

To see how this approach works, recall that the solution to the finite horizon problem for periods $0, \ldots, (\tau - 1)$ can be viewed as a two point boundary value problem where the initial state vector y_0 is set to some arbitrary vector y and the costate vector at the terminal date μ_{τ} is set to zero. Suppose for simplicity that A_{yy} is nonsingular. By iterating on relation (3.3.14), we find that

$$\hat{M} \begin{bmatrix} y_{\tau} \\ 0 \end{bmatrix} = \begin{bmatrix} y_0 \\ \mu_0 \end{bmatrix}, \qquad (3.4.10)$$

where

$$\hat{M} \equiv M^{-\tau}$$

To approximate the matrix P_y , we solve (3.4.10) for the initial costate vector μ_0 as a function of y_0 . Partitioning \hat{M} conformably to the state-costate partition, we see that

$$\hat{M}_{11}y_{\tau} = y_0, \qquad \hat{M}_{21}y_{\tau} = \mu_0.$$

Therefore, the implicit initialization of the costate vector is

$$\mu_0 = \hat{M}_{21} \left(\hat{M}_{11} \right)^{-1} y_0,$$

and our approximation for the matrix P_y is given by $\hat{M}_{21}(\hat{M}_{11})^{-1}$.

What is needed to implement this approach is a way to compute \hat{M} when the horizon τ is large. Expanding the horizon one period at a time corresponds to multiplying the matrix M^{-1} , τ times in succession. However, when τ is chosen to be a power of two, computations can be sped up by using

$$M^{-2^{k+1}} = \left(M^{-2^k}\right) M^{-2^k}.$$
(3.4.11)

As a consequence, when $\tau = 2^{j}$, the desired matrix can be computed in j iterations instead of 2^{j} iterations, which explains the name *doubling algorithm*.

Given that the matrix M^{-1} has unstable eigenvalues, direct iterations on (3.4.11) can be very unreliable. Clearly, the sequence of matrices $\{M^{-2^k}\}$ diverges. One of the features of a doubling algorithm is to transform these computations into matrix iterations that converge. Another feature is that a doubling algorithm exploits the fact that the matrix M is symplectic. Symplectic matrices have several nice properties.¹¹ We have already seen that their eigenvalues come in reciprocal pairs. In addition, the product of symplectic matrices is symplectic, and the inverse of a symplectic matrix is symplectic. Moreover, for any symplectic matrix S, the matrices $S_{21}(S_{11})^{-1}$ and $(S_{11})^{-1}S_{12}$ are both symmetric and

$$S_{22} = (S'_{11})^{-1} + S_{21} (S_{11})^{-1} S_{12}$$

= $(S'_{11})^{-1} + S_{21} (S_{11})^{-1} S_{11} (S_{11})^{-1} S_{12}.$

Therefore, a $(2n \times 2n)$ symplectic matrix can be represented in terms of the three $n \times n$ matrices $\alpha = (S_{11})^{-1}$, $\beta = (S_{11})^{-1}S_{12}$, $\gamma = S_{21}(S_{11})^{-1}$, the latter two of which are symmetric.

The doubling algorithm described by Anderson (1978) and Anderson and Moore (1979) exploits such a representation by using the following parameterization of M^{-2^k} :

$$M^{-2^{k}} = \begin{bmatrix} (\alpha_{k})^{-1} & (\alpha_{k})^{-1} \beta_{k} \\ \gamma_{k} (\alpha_{k})^{-1} & \alpha'_{k} + \gamma_{k} (\alpha_{k})^{-1} \beta_{k} \end{bmatrix},$$

where the $n \times n$ matrices $\alpha_k, \beta_k, \gamma_k$ are given by the recursions

$$\alpha_{k+1} = \alpha_k \left(I + \beta_k \gamma_k\right)^{-1} \alpha_k$$

$$\beta_{k+1} = \beta_k + \alpha_k \left(I + \beta_k \gamma_k\right)^{-1} \beta_k \alpha'_k$$

$$\gamma_{k+1} = \gamma_k + \alpha'_k \gamma_k \left(I + \beta_k \gamma_k\right)^{-1} \alpha_k.$$

(3.4.12)

While this alternative parameterization introduces a matrix inverse into the recursions (3.4.12) that is absent in (3.4.11), the matrix $I + \beta_k \gamma_k$ being inverted is only *n* dimensional. The nonsingularity of this matrix for all *k* is established in Kimura (1988). To initialize the doubling algorithm, we simply deduce the implicit parameterization of M^{-1} given in partitioned form by

$$M^{-1} = N^{-1}L = \begin{bmatrix} A_{yy}^{-1} & A_{yy}^{-1}B_{y}R^{-1}B_{y}' \\ Q_{yy}A_{yy}^{-1} & Q_{yy}A_{yy}^{-1}B_{y}R^{-1}B_{y}' + A_{yy}' \end{bmatrix}, \quad (3.4.13)$$

¹¹ There is a variation of the Schur algorithm that exploits the symplectic structure of M. See pages 431-434 of Petkov *et al.* (1991) for an overview of this algorithm.

which leads to the initializations

$$\alpha_0 = A_{yy}, \quad \beta_0 = B_y R^{-1} B_y', \quad \gamma_0 = Q_{yy}.$$

While our derivation took the matrix A_{yy} to be nonsingular, Anderson (1978) argues that the doubling algorithm is more generally applicable.

A convenient feature of this parameterization is that there are known conditions under which the matrix sequences $\{\alpha_k\}, \{\beta_k\}, \{\gamma_k\}$ converge. When the pair (A_{yy}, D_y) is detectable, then the sequence $\{\gamma_k\}$ is nondecreasing and converges to the matrix P_y . (Here we are adopting the usual partial ordering for positive semidefinite matrices.) As noted by Kimura (1988, Theorem 5), under the same restrictions, the sequence $\{\beta_k\}$ is nondecreasing and converges to a positive semidefinite matrix P_y^* associated with a "dual" to the *deterministic* regulator problem.

The convergence of the $\{\alpha_k\}$ sequence is more problematic. Unfortunately, without simultaneous convergence of $\{\alpha_k\}$, it is not evident that iterations of the form given in (3.4.12) can be used as the basis of a numerical algorithm. If this latter sequence diverges, small numerical errors may get magnified, causing the resulting algorithm to be poorly behaved. Kimura (1988) provides some sufficient conditions for $\{\alpha_k\}$ to converge to a matrix of zeros. His sufficient conditions are used to guarantee that either P_y or P_y^* is nonsingular.

As we noted previously, a sufficient condition for P_y to be nonsingular is that the pair (A_{yy}, D_y) be observable. Sufficient conditions for the nonsingularity of the matrix P_y^* are that (i) (A_{yy}, B_y) is controllable; and (ii) (A_{yy}, D_y) is detectable (Kimura 1988). Recall that controllability is often achieved by our *a priori* partitioning of the state vector into *endogenous* and *exogenous* components. Thus for our purposes, the restrictions guaranteeing the nonsingularity of P_y^* may be of particular interest. Even so, detectability is too strong for some of our applications.

To apply a doubling algorithm more generally, we sometimes modify the control problem by adding small quadratic penalties to linear combinations of the states and controls. As long as these penalties are sufficient to guarantee that either P_y or P_y^* is nonsingular, we are assured of convergence of all three sequences. Of course, there is a danger that the penalty distorts the solution to the original control problem in a nontrivial way, which must be checked in practice.

3.4.3.1. Initialization from a positive definite matrix

Instead of adding small quadratic penalties to the objective function for each calendar date, we could add a terminal penalty to the finite horizon approximation to the control problem. From Chan, Goodwin and Sin (1984), it is known that iterations on the Riccati difference equation converge to the unique stabilizing solution whenever the Riccati equation is initialized at a positive definite matrix.¹² Initializing the Riccati difference equation at a positive definite matrix is equivalent to imposing a terminal penalty that is a negative definite quadratic form in the state vector. We will now show how to initialize the doubling algorithm to impose a terminal penalty. This will permit us to compute P_{y} via a doubling algorithm for a richer class of control problems.

Consider first a finite time horizon problem with a quadratic penalty on the terminal state. We select this penalty so that the terminal multiplier $\mu_{\tau} = P_o y_{\tau}$ for some positive definite matrix P_o . Then equation (3.4.10) is altered to be

$$\hat{M} \begin{bmatrix} I \\ P_o \end{bmatrix} y_{\tau} = \begin{bmatrix} y_0 \\ \mu_0 \end{bmatrix}.$$
(3.4.14)

Build a matrix K

$$K \equiv \begin{bmatrix} I & 0 \\ P_o & I \end{bmatrix}.$$

Then equation (3.4.14) can be rewritten as

$$K^{-1}\hat{M}KK^{-1}\begin{bmatrix}I\\P_o\end{bmatrix}y_{\tau} = K^{-1}\begin{bmatrix}y_0\\\mu_0\end{bmatrix}.$$

Equivalently,

$$M^* \begin{bmatrix} y_\tau \\ 0 \end{bmatrix} = \begin{bmatrix} y_0 \\ \mu_0 - P_o y_0 \end{bmatrix},$$

where

$$M^* = K^{-1}\hat{M}K.$$

Partitioning M^* consistently with the state-costate vector, the implicit initialization of the costate vector is now

$$\mu_0 = P_o y_0 + M_{12}^* \left(M_{11}^* \right)^{-1} y_0,$$

¹² Here we are using the fact that the pair (A_{yy}, B_y) is stabilizable and that there exists a solution to the *deterministic regulator problem* when constraint (3.2.1) is imposed. The result follows from (i) and (iii) of Theorem 3.1 and Theorem 4.2 of Chan, Goodwin and Sin (1984).

and our approximation for P_y is given by $M_{12}^*(M_{11}^*)^{-1} + P_o$.

We are now left with computing the matrix M^* when the horizon τ is very large. Notice that

$$M^* = \left(K^{-1}MK\right)^{-\tau}.$$

It is straightforward to verify that because M is symplectic, so is $K^{-1}MK$. This means that doubling algorithm (3.4.12) is applicable for computing $(K^{-1}MK)^{-2^k}$; however, the initializations must be altered. The new initializations can be deduced by looking at the implicit parameterization of the symplectic matrix $K^{-1}M^{-1}K$, and they are given by

$$\alpha_{0} = (I + B_{y}R^{-1}B_{y}'P_{o})^{-1}A_{yy}$$

$$\beta_{0} = (I + B_{y}R^{-1}B_{y}'P_{o})^{-1}B_{y}R^{-1}B_{y}'$$

$$\gamma_{0} = Q_{yy} - P_{o} + A_{yy}'P_{o}(I + B_{y}R^{-1}B_{y}'P_{o})^{-1}A_{yy}.$$

(3.4.15)

Not surprisingly, the original initializations coincide with setting P_o to zero in (3.4.15).

There are two related advantages to these initializations over the previous ones. First, the sequence $\{\gamma_j\}$ converges to $P_y - P_o$ whenever P_o is positive definite. This follows from the Riccati difference equation convergence described previously and does not require that (A_{yy}, D_y) be detectable. Second, the sequence $\{\beta_j\}$ converges and satisfies the bounds

$$0 \le \beta_j \le \left(P_o\right)^{-1}$$

even when (A_{yy}, D_y) is not detectable.¹³ Although we do not have a complete characterization of convergence of the resulting algorithm, all three matrix

$$\beta_j = \left(I + \beta_j^* P_o\right)^{-1} \beta_j^*.$$

Exploiting the nonsingularity of P_o , the following equivalent formula can be deduced:

$$\beta_j = (P_o)^{-1} - (P_o + P_o \beta_j^* P_o)^{-1}$$

¹³ The convergence and bound can be established as follows. Let $\{\beta_j^*\}$ denote the sequence starting from the original initialization. Then it is straightforward to show that

The reported bound follows immediately. The sequence $\{\beta_j^*\}$ is monotone increasing because it is a subsequence of Riccati difference equation iterations for a dual problem initialized at zero. Therefore, the sequence $\{\beta_j\}$ is also monotone increasing. Given the upper bound $(P_o)^{-1}$, this latter sequence must converge.

sequences (including $\{\alpha_j\}$) are guaranteed to converge with these alternative initializations if they converge with the original ones.

In summary, the steps for implementing the doubling algorithm are

- (1) initialize α_0 , β_0 , and γ_0 according to (3.4.15);
- (2) iterate in accordance with (3.4.12);
- (3) form P_y as the limit of $\{\gamma_k\} + P_o$.

3.4.3.2. Application to continuous time

As noted by Anderson (1978) and Kimura (1989), a doubling algorithm for a discrete-time symplectic system can be used to solve a continuous-time Hamiltonian system. Recall that in our discussion of solving control problems via a matrix sign algorithm, we showed how to covert a discrete-time symplectic system into a continuous-time Hamiltonian system. To apply a doubling algorithm, we want to "invert" this mapping, *e.g.*, given a Hamiltonian matrix H, we construct a symplectic pencil with the same stable deflating subspace. The symplectic pencil associated with H is given by $\lambda(I + H) - (I - H)$. By adopting a very similar argument as before, we found it easy to show that the generalized eigenvectors for the constructed pencil coincide with the eigenvectors of the original Hamiltonian matrix H. Moreover, the classification of stable and unstable (generalized) eigenvalues is preserved.

3.4.4. Matrix sign algorithm

In section 3.3.3 we showed how to compute P_y from the sign of the Hamiltonian matrix for a continuous-time state-costate system. To compute P_y for a symplectic pencil $\lambda L - N$, we first form the Hamiltonian matrix

$$H = (L - N)^{-1} (L + N)$$

and then compute sign(H). For this to be a viable solution method, we must be able to compute sign(H) easily.

There are alternative matrix sign algorithms. An algorithm advocated by Roberts (1980) and Denman and Beavers (1976) is to average a matrix and its inverse:

$$G_0 = H$$

$$G_{k+1} = G_k + (1/2) \left[(G_k)^{-1} - G_k \right], k = 0, 1, \dots$$
(3.4.16)

To speed up convergence, Gardiner and Laub (1986) suggest using the recursion

$$G_0 = H, \quad G_{k+1} = (1/2\epsilon_k) \left(G_k + \epsilon_k^2 G_k^{-1}\right),$$

where

$$\epsilon_k = \left|\det G_k\right|^{1/n}.\tag{3.4.17}$$

Bierman (1984) and Byers (1987) propose a further refinement, which exploits the fact that the matrix G_k is a Hamiltonian matrix for each k. Recall that if H is a Hamiltonian matrix, then JH is symmetric where

$$J = \begin{bmatrix} 0 & -I \\ I & 0 \end{bmatrix}.$$

Hence

$$JG_{k+1} = \frac{1}{2\epsilon_k} \left(JG_k + {\epsilon_k}^2 J J G_k^{-1} J \right), \qquad (3.4.18)$$

where ϵ_k is either set to one as in the original sign algorithm or set via formula (3.4.17) using JG_k in place of G_k . Consequently, it suffices to compute the sequence of symmetric matrices $\{JG_k\}$ recursively via (3.4.18) starting from the initialization JH.¹⁴

In summary, the steps for implementing a matrix sign algorithm are

- (1) form the matrices L and N in (3.3.4);
- (2) compute the sign of $G = (L N)^{-1}(L + N);$
- (3) compute P_y by solving the over-determined system

$$\begin{bmatrix} G_{12} \\ G_{22} + I \end{bmatrix} P_y = -\begin{bmatrix} G_{11} + I \\ G_{21} \end{bmatrix}$$
(3.4.19)

for P_y .

As noted in Anderson (1978), the original sign algorithm (3.4.16) also can be viewed as a doubling algorithm. Interpreted in this manner, it uses (at least

 $^{^{14}\,}$ Kenney, Laub and Papadopoulos (1993) and Lu and Lin(1993) discuss further improvements to the matrix sign algorithm.

implicitly) an alternative parameterization of the symplectic matrix M^{-1} to that used in the doubling algorithm (3.4.12). Both recursions entail inverting a matrix. While recursion (3.4.18) requires that a symmetric $(2n \times 2n)$ matrix be inverted in each iteration, the doubling algorithm (3.4.12) requires that a nonsymmetric $n \times n$ matrix be computed at each iteration.

3.5. Solving the augmented regulator problem

So far, we have shown how to compute the matrix F_y , which provides us with the optimal control law for the *deterministic regulator problem*. This matrix also gives us a piece of the solution to the *augmented control problem* and, hence, to the problem of interest: the *discounted stochastic regulator problem*. The missing ingredient is the matrix F_z , where the optimal control law for the *augmented regulator problem* is given by $v_t = -F_y y_t - F_z z_t$. In this section, we show that F_z can be calculated by solving a particular Sylvester equation.

We start by forming a Lagrangian modified to incorporate the exogenous state vector sequence $\{z_t\}$:

$$\mathcal{L} = -\sum_{t=0}^{\infty} \left[y_t' Q_{yy} y_t + 2y_t' Q_{yz} z_t + v_t' R v_t + 2\mu_{t+1}' \left(A_{yy} y_t + A_{yz} z_t + B_y v_t - y_{t+1} \right) \right],$$

where the evolution of the forcing sequence is given by

$$z_{t+1} = A_{zz} z_t. (3.5.1)$$

First-order necessary conditions for the maximization of \mathcal{L} with respect to $\{v_t\}_{t=0}^{\infty}$ and $\{y_t\}_{t=0}^{\infty}$ are

$$v_t: \qquad Rv_t + B_y' \mu_{t+1} = 0, \qquad t \ge 0 \tag{3.5.2}$$

$$y_t: \qquad \mu_t = Q_{yy}y_t + Q_{yz}z_t + A_{yy}'\mu_{t+1}, \qquad t \ge 0.$$
(3.5.3)

Solve equation (3.5.2) for v_t ; substitute it into the state equation; and stack the resulting equation along with (3.5.3) and (3.5.1) as composite system

$$L^{a} \begin{bmatrix} y_{t+1} \\ \mu_{t+1} \\ z_{t+1} \end{bmatrix} = N^{a} \begin{bmatrix} y_{t} \\ \mu_{t} \\ z_{t} \end{bmatrix},$$

where

$$L^{a} \equiv \begin{bmatrix} I & B_{y}R^{-1}B'_{y} & 0\\ 0 & A_{yy}' & 0\\ 0 & 0 & I \end{bmatrix}, \qquad N^{a} \equiv \begin{bmatrix} A_{yy} & 0 & A_{yz}\\ -Q_{yy} & I & -Q_{yz}\\ 0 & 0 & A_{zz} \end{bmatrix}.$$
(3.5.4)

As with the *deterministic regulator problem*, the relevant solution is the one that stabilizes the state-costate vector for any initialization of y_0 and z_0 . Hence we seek a characterization of the multiplier μ_t of the form

$$\mu_t = P \begin{bmatrix} y_t \\ z_t \end{bmatrix},$$

such that the resulting composite sequence $[y_t' \ \mu_t' \ z_t']'$ is in the stable deflating subspace of the augmented pencil $\lambda L^a - N^a$. Assuming for the moment that a solution P exists, it must be the case that $P = [P_y \ P_z]$, where P_y is the Riccati equation solution that was characterized in section 3.3, and P_z is a matrix that has not yet been characterized. To see why this must be the case, note that the solution to the *augmented regulator problem* with $z_0 = 0$ coincides with the solution to the *deterministic regulator problem*. We have previously shown that P_y is a matrix, such that all vectors in the deflating subspace of the pencil $\lambda L - N$ can be represented as $[y' \ y'P_y]'$. When the forcing sequence is initialized at zero, so that it remains there for all t, it must also be the case that $[y' \ y'P_y \ 0]'$ is in the stable deflating subspace of the augmented pencil $\lambda L^a - N^a$. This justifies our previous claim that the solution to the *deterministic regulator problem* is a piece of the solution to the *augmented regulator problem*.

To deduce the control law associated with the matrix P, we substitute P into (3.5.4), which yields

$$L^{a} \begin{bmatrix} y_{t+1} \\ P_{y}y_{t+1} + P_{z}z_{t+1} \\ z_{t+1} \end{bmatrix} = N^{a} \begin{bmatrix} y_{t} \\ P_{y}y_{t} + P_{z}z_{t} \\ z_{t} \end{bmatrix}.$$

Write the three equations in this composite system separately:

$$(I + B_y R^{-1} B_y' P_y) y_{t+1} + B_y R^{-1} B_y' P_z z_{t+1} = A_{yy} y_t + A_{yz} z_t A_{yy}' P_y y_{t+1} + A_{yy}' P_z z_{t+1} = (P_y - Q_{yy}) y_t + (P_z - Q_{yz}) z_t z_{t+1} = A_{zz} z_t.$$

$$(3.5.5)$$

Substitute the last equation into the first and solve for y_{t+1} :

$$y_{t+1} = \left(I + B_y R^{-1} B_y' P_y\right)^{-1} \left[A_{yy} y_t + \left(A_{yz} - B_y R^{-1} B_y' P_z A_{zz}\right) z_t\right].$$

It follows from relation (3.3.9) that this evolution equation for y_t can be rewritten as

$$y_{t+1} = (A_{yy} - B_y F_y) y_t + (A_{yz} - B_y F_z) z_t, \qquad (3.5.6)$$

where F_y and F_z are given by

$$F_{y} \equiv \left(R + B_{y}' P_{y} B_{y}\right)^{-1} B_{y}' P_{y} A_{yy},$$

$$F_{z} \equiv \left(R + B_{y}' P_{y} B_{y}\right)^{-1} B_{y}' \left(P_{y} A_{yz} + P_{z} A_{zz}\right).$$
(3.5.7)

For the reasons given previously, our construction of F_y coincides with (3.3.11) used to represent the optimal control law for the *deterministic regulator problem*. Stability of the state vector sequence $\{y_t\}$ is guaranteed by evolution equation (3.5.6) because the matrix $A_{yy} - B_y F_y$ is the same matrix that appears in the state evolution equation for the *deterministic regulator problem* under the optimal control law. Since the solution to the *deterministic regulator problem* is stable by design, the eigenvalues of $A_{yy} - B_y F_y$ have absolute values that are strictly less than one. The optimal control law for the *augmented regulator problem* is given by

$$w_t = -F_y y_t - F_z z_t.$$

The matrix F_z can be computed using formula (3.5.7) once we know P_z . We now show that P_z is the solution to a *Sylvester equation*. Premultiply (3.5.6) by $A_{yy}'P_y$:

$$A_{yy}'P_{y}y_{t+1} = A_{yy}'P_{y}\left(A_{yy} - B_{y}F_{y}\right)y_{t} + A_{yy}'P_{y}\left(A_{yz} - B_{y}F_{z}\right)z_{t}.$$
 (3.5.8)

Using formula (3.5.7), we rewrite the coefficient matrix on z_t as

$$A_{yy}'P_y(A_{yz} - B_yF_z) = (A_{yy} - B_yF_y)'(P_yA_{yz} + P_zA_{zz}) - A_{yy}'P_zA_{zz}.$$

To obtain an alternative formula for this coefficient, substitute the last equation of (3.5.5) into the second equation and solve for $A_{yy}'P_yy_{t+1}$:

$$A_{yy}'P_{y}y_{t+1} = (P_z - Q_{yz} - A_{yy}'P_zA_{zz})z_t + (P_y - Q_{yy})y_t.$$
 (3.5.9)

Equating coefficients on z_t in (3.5.8) and (3.5.9) results in

$$(A_{yy} - B_y F_y)' (P_y A_{yz} + P_z A_{zz}) - A_{yy}' P_z A_{zz} = P_z - Q_{yz} - A_{yy}' P_z A_{zz}.$$

Rewriting this in the form of a Sylvester equation (in the unknown matrix P_z), we have that

$$P_{z} = Q_{yz} + (A_{yy} - B_{y}F_{y})' P_{y}A_{yz} + (A_{yy} - B_{y}F_{y})' P_{z}A_{zz}.$$
 (3.5.10)

As already noted, the matrix $(A_{yy} - B_y F_y)$ has only stable eigenvalues. Also, we assumed that the matrix A_{zz} has only stable eigenvalues (Assumption 4). These restrictions are sufficient for there to exist a unique solution P_z to (3.5.10). Up to now, our discussion proceeded under the presumption that there exists a matrix P, such that by setting $\mu_t = P \begin{bmatrix} y_t \\ z_t \end{bmatrix}$, we stabilize the state vector sequence. We can now work backwards using the (unique) solution to the Sylvester equation to show that indeed such a matrix P does exist.

3.6. Computational techniques for solving Sylvester equations

A Sylvester equation is represented by

$$M = W + SMT, (3.6.1)$$

where the matrices W, S, and T are specified in advance and M is the matrix to be computed. Consistent with (3.5.10), the matrices S and T have stable eigenvalues.¹⁵ There is a variety of ways to depict the solution to a Sylvester equation. One is to vectorize (3.6.1) as

$$[I - T' \otimes S] \operatorname{vec}(M) = \operatorname{vec}(W), \qquad (3.6.2)$$

where $\operatorname{vec}(\cdot)$ denotes stacks of the columns of a matrix argument. (To derive (3.6.2) from (3.6.1), use the identity $\operatorname{vec}(SMT) = [T' \otimes S]\operatorname{vec}(M)$). Hence

 $^{^{15}}$ We have recycled some of the notation used in previous sections.

 $\operatorname{vec}(M)$ is the solution to a linear equation system. Alternatively, M is given by the infinite sum

$$M = \sum_{j=0}^{\infty} S^{j} W T^{j}.$$
 (3.6.3)

This representation can be deduced by iterating on equation (3.6.1), starting from any initial matrix with the appropriate dimensions.

We consider two types of algorithms for computing M:

(i) Hessenburg-Schur algorithm;

(ii) doubling algorithm.

The Hessenberg-Schur algorithm uses a Schur decomposition of the matrix T to convert a single Sylvester equation to a collection of much smaller Sylvester equations, each of which can be vectorized as in (3.6.2). A Hessenberg decomposition of the matrix S is used further to simplify the calculations. The doubling algorithm is an iterative algorithm that approximates the infinite sum on the right-hand side of (3.6.3) by a finite sum. As with the doubling algorithm for solving a Riccati equation, the number of terms included in the finite sum approximation "doubles" at each iteration.

3.6.1. The Hessenberg-Schur algorithm

As suggested by Bartels and Stewart (1972), one strategy for solving Sylvester equations entails block triangularizing the matrices T and/or S. We follow Golub, Nash and Van Loan (1979) by forming a Schur decomposition of the matrix $T: V'TV = \hat{T}$, where V is an orthogonal matrix and \hat{T} is upper block triangular with row and column blocks that are either one or two dimensional (see section 3.4.1 for a formal definition). Postmultiply Sylvester equation (3.6.1) by V and rewrite the equation as

$$\hat{M} = \hat{W} + \hat{S}\hat{M}\hat{T}, \qquad (3.6.4)$$

where $\hat{M} = MV$, $\hat{W} = WV$, and $\hat{S} = S$. Notice that (3.6.4) is in the form of a Sylvester equation in the matrix \hat{M} .

The block triangularity of \hat{T} can now be exploited to reduce (3.6.4) into m smaller Sylvester equations, where m is the number of row and column blocks

of \hat{T} . Write the matrix \hat{T} in partitioned form as

$$\hat{T} = \begin{bmatrix} \hat{T}_{11} & \hat{T}_{12} & \dots & \hat{T}_{1m} \\ 0 & \hat{T}_{22} & \dots & \hat{T}_{2m} \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & \hat{T}_{mm} \end{bmatrix}$$

Use the column partition of W to partition \hat{M} and \hat{W} , and let \hat{M}_j and \hat{W}_j denote the corresponding j^{th} partitions. Decompose Sylvester equation (3.6.4):

$$\hat{M}_1 = \hat{W}_1 + \hat{S}\hat{M}_1\hat{T}_{11} \tag{3.6.5}$$

$$\hat{M}_j = \hat{W}_j + \hat{S} \sum_{k=1}^{j-1} \hat{M}_k \hat{T}_{kj} + \hat{S} \hat{M}_j \hat{T}_{jj}, \quad j = 2, ..., m.$$
(3.6.6)

Notice that (3.6.5) is a Sylvester equation in \hat{M}_1 and that (3.6.6) is a Sylvester equation in \hat{M}_j as long as the matrices \hat{M}_k for k = 1, 2, ..., j - 1 have already been computed. Thus these *m* Sylvester equations can be solved sequentially as linear equations using vectorization (3.6.2).

An additional refinement advocated by Golub, Nash and Van Loan (1979) entails taking a Hessenberg decomposition of the matrix S.¹⁶

Definition: The Hessenberg decomposition of the square matrix S is an orthogonal matrix U and a matrix \hat{S} that has all zeros below the first subdiagonal, such that $S = U\hat{S}U'$.

In addition to postmultiplying equation (3.6.1) by V, we now also premultiply this equation by U'. Equation (3.6.4) continues to hold with $\hat{M} = U'MV$, $\hat{W} = U'WV$, and $\hat{S} = U'SU$. This Sylvester equation can still be decomposed as in (3.6.5) and (3.6.6). With \hat{S} in Hessenberg form, we can solve these latter Sylvester equations more efficiently using an equation solver designed for Hessenberg systems.¹⁷

In summary, the steps for implementing a Hessenberg-Schur algorithm for computing ${\cal P}_z$ are

 $^{^{16}\,}$ Alternatively, we could take the Schur decomposition of S as proposed by Bartels and Stewart (1972).

¹⁷ Interesting variations on the Hessenberg-Schur algorithm have been proposed by Hammarling(1982) and Gardiner *et al.* (1992).

- (i) form the matrices $W = Q_{yz} + (A_{yy} B_y F_y)' P_y A_{yz}$, $S = (A_{yy} B_y F_y)'$, and $T = A_{zz}$;
- (ii) form a Hessenberg decomposition $S = U\hat{S}U'$ and a Schur decomposition $T = V\hat{T}V'$;
- (iii) compute the solution \hat{M} to (3.6.5) and (3.6.6) and form $P_z = U\hat{M}V'$.

Since the Hessenberg decomposition of a matrix can be computed faster than the real Schur decomposition, one should always arrange the Sylvester equation so that the Hessenberg decomposition is taken of the matrix $(A_{yy} - B_y F_y)'$ or A_{zz} , whichever has more entries. The steps just described should be implemented if there are more elements in the vector y_t than z_t . If z_t has more elements, then the alternative Sylvester equation

$$P_{z}' = Q_{yz}' + A'_{yz}P_{y}(A_{yy} - B_{y}F_{y}) + A'_{zz}P_{z}'(A_{yy} - B_{y}F_{y})'$$

should be solved for the matrix P_z' .¹⁸

3.6.2. Doubling algorithm

The doubling algorithm for Sylvester equations iterates on

$$\alpha_{k+1} = \alpha_k \alpha_k$$

$$\beta_{k+1} = \beta_k \beta_k$$

$$\gamma_{k+1} = \gamma_k + \alpha_k \gamma_k \beta_k$$

(3.6.7)

to convergence, where $\alpha_0 = S$, $\beta_0 = T$, and $\gamma_0 = W$. By repeated substitution, it can be shown that

$$\gamma_k = \sum_{j=0}^{2^k - 1} S^j W T^j.$$

In other words, each iteration doubles the number of terms in the sum.¹⁹

¹⁸ In numerical work in Anderson, Hansen, McGrattan, and Sargent (1996), we formed the Hessenberg decomposition of a matrix using MATLAB subroutine HESS and the Schur decomposition of a matrix with SCHUR. We solved Hessenberg systems using the routines HSFA and HSSL, which are part of the package described in Gardiner *et al.* (1992). See pages 364-370 of Golub and Van Loan (1989) for how to compute the Hessenberg decomposition.

¹⁹ This algorithm is a slight generalization of the doubling algorithm for Lyapunov equations discussed in Anderson and Moore (1979). A Lyapunov equation is a Sylvester equation in which S = T'.
To use this doubling algorithm to compute P_z

- (i) initialize $\alpha_0 = (A_{yy} B_y F_y)'$, $\beta_0 = A_{zz}$, and $\gamma_0 = Q_{yz} + (A_{yy} B_y F_y)' P_y A_{yz}$;
- (ii) iterate in accordance to (3.6.7);
- (iii) form P_z as the limit of $\{\gamma_k\}$.

3.7. Concluding remarks

This chapter has focused on computational details for the optimal linear regulator. Many aspects of these calculations will recur in various settings below. Indeed, key ideas and formulas in all of the subsequent chapters of this book build directly or indirectly on results in this chapter. Thus, in chapter 4, we see how the Kalman filter emerges as the dual of the optimal linear regulator. Chapter 6 uses invariant subspace methods to prove the equivalence of alternative ways of formulating a robust control problem. Chapter 16 uses a Lagrangian formulation and invariant subspace methods to construct robust decision rules for controlling forward looking models. As already indicated in chapter 2, the optimal linear regulator can be induced to do all of the hard work in computing a robust rule for such models.

Chapter 4 The Kalman filter

 ... we are always searching for something hidden or merely potential or hypothetical, following its traces whenever they appear on the surface.
 Italo Calvino, Six Memos for the Next Millenium, 1996, p. 77.

4.1. Introduction

The Kalman filter is a recursive method for computing linear least squares estimates of sequences of random vectors comprising hidden states and future observables. The states and observables are described by a *known* linear statespace system that is perturbed by Gaussian shocks with zero mean and known covariances.

Remarkably, the Kalman filter formulas are identical with those for an optimal linear regulator, a fact that reflects the duality of filtering and control, the subject of this chapter. Following Whittle (1990, 1996), we formulate a filtering problem in terms of a Lagrangian. After performing minimizations and maximizations in a particular order, an optimal linear regulator problem emerges with the flow of time reversed. We therefore say that the linear regulator problem is *dual* to the Kalman filter, and *vice versa*.

The Kalman filter is a powerful tool in economics and econometrics because it accomplishes many tasks, including these: (1) it efficiently computes the Wold and autoregressive representations associated with an economic model whose equilibrium can be represented as a linear state space system;¹ (2) By recovering an autoregressive representation, it enables computing the likelihood function of a linear model recursively; (3) by building upon (2), it can be used to infer the econometric implications of aggregation over time; and (4) it is the basic tool for estimating and forecasting hidden factors in linear models. Items (1)–(4) make the Kalman filter an essential tool in deducing the observable implications for an

- 99 -

¹ A common practice in the real business cycle literature is to approximate an equilibrium as a linear state space system in logarithms of state variables. That enables the application of the Kalman filter to obtain the vector autoregressive representation and the likelihood. , For examples, see Schorfveide (2000XXX) and Otrok (2001XXX).

important class of models whose equilibria occur, or can be well approximated, in the form of a linear state space system. 2

Before getting into the details, we first state the Kalman filtering problem and its solution, then assert the associated optimal linear regulator problem for which it is the dual. The remaining sections of the chapter fill in the details required to prove the duality of the filtering and control problems.

We assume throughout this chapter that the state-space model is true, so that issues of model approximation are not in play. Chapters 13 and 14 will formulate filtering problems in setings where the decision maker suspects model misspecification and therefore wants a robust filter.

4.2. Review of Kalman filter and preview of main result

Throughout this chapter, we let x_t denote a state vector at time t and y_t a vector of possibly noise-ridden observations on linear combinations of x_{t-1} . This section uses a convention for indexing time that differs from the one used in the remainder of the chapter. We temporarily use this timing convention because we shall use it again in chapter 13 and because it leads to a dual control problem in which the direction of time matches the one we used in chapters 2 and 3. To attain that familiar representation for the control problem, for the filtering problem we have to let larger indexes t recede further into the past. We begin with a simple and famous example.

 $^{^2}$ So far as first and second moments are concerned, those implications are characterized by a vector autoregression. Using the Kalman filter is the easiest way to obtain the autoregressive representation. See Hansen and Sargent (200XXX, chapter 8).

4.2.1. Muth's problem as an example

John F. Muth (1960) applied classical filtering methods to discover a stochastic process for income for which Milton Friedman's (1956XX) adaptive expectations scheme would be an optimal estimator of permanent income. Muth's problem can be formulated recursively using the Kalman filter. Where x_{-t} is a scalar state variable and y_{-t} is a scalar observed variable at time $-t, t \ge 0$, consider the state space system

$$x_{-t} = ax_{-t-1} + \begin{bmatrix} c & 0 \end{bmatrix} \epsilon_{-t}$$
(4.2.1*a*)

$$y_{-t} = gx_{-t-1} + \begin{bmatrix} 0 & d \end{bmatrix} \epsilon_{-t}$$
(4.2.1b)

where a, g, c, d are scalars and ϵ_{-t} is an i.i.d. (2×1) vector of Gaussian random variables with mean zero and covariance matrix I. To analyze Milton Friedman's concept of permanent income, Muth set a = 1, g = 1 and c > 0, d > 0. He regarded x_{-t} as a permanent component of income and $d\epsilon_{2,-t}$ as transitory income, while y_{-t} is measured income at -t. A consumer facing an income process with this structure wants to estimate his permanent income. Thus, he wants to compute $\hat{x}_{-t} \equiv E[x_{-t}|y^{-t}]$ where y^{-t} denotes the infinite history of $[y_{-t}, y_{-t-1}, \ldots]$. That is, the consumer wants to form an estimator \hat{x}_{-t} that is a measurable function of the infinite history y^{-t} and that minimizes $E[(x_{-t} - \hat{x}_{-t})^2|y^{-t}]$.

The Kalman filter attains Muth's solution of this problem.³ The solution for the optimal estimator takes the recursive form $\hat{x}_{-t} = (a - Kg)\hat{x}_{-t-1} + Ky_{-t}\hat{x}_{-t}$, which can also be represented as

$$\hat{x}_{-t} = K \sum_{j=0}^{\infty} \left(a - Kg\right)^j y_{-t-j}$$
(4.2.2)

where K is the Kalman gain. Equation (4.2.2) expresses the consumer's estimate of the permanent component of his income as a geometric weighted sum of past income levels. The conditional variance of this estimator is $\Sigma = E \left[x_{-t} - \hat{x}_{-t} | y^{-t} \right]^2$. The Kalman filter gives a way to compute Σ and K.

³ Muth solved the problem using classical (i.e., non-recursive, methods.

4.2.2. The dual to Muth's filtering problem

The dual to Muth's filtering problem is the optimal linear regulator

$$-\Sigma\lambda_0^2 \equiv \max_{\{\mu_t\}} -\sum_{t=0}^{\infty} \left(c^2\lambda_t^2 + d^2\mu_t^2\right)$$
(4.2.3)

where the maximization is subject to the law of motion

$$\lambda_{t+1} = a\lambda_t + g\mu_t, \tag{4.2.4}$$

with λ_0 given, and where a, g, c, d take the same values as in Muth's problem. Problem (4.2.3), (4.2.4) has a solution in the form of a feedback rule

$$\mu_t = -K\lambda_t \tag{4.2.5}$$

where K is the same scalar that emerges from the Kalman filter, and the matrix Σ in the value function $-\Sigma\lambda_0^2$ is the state covariance matrix that emerges from the Kalman filter. In this chapter, we shall interpret the λ 's as Lagrange multipliers associated with the Kalman filtering problem.

For particular values of a, g, c, d, we invite the reader to use the Matlab program olrp.m to solve the regulator problem and kfilter.m to solve the Kalman filtering problem, and thereby to verify numerically the duality that we have asserted. In the next section, we verify duality analytically and in the process tell why the adjective 'dual' is appropriate, in the sense of mathematical programming. is appropriate. But first we state a more general versions of the Kalman filter problem and the associated dual optimal linear regulator problem.

4.2.3. The filtering problem

Consider the following optimal filtering problem that generalizes Muth's problem. For $t \ge 0$, a state vector x_{-t} and an observation vector y_{-t} satisfy⁴

$$x_{-t} = Ax_{-t-1} + C\epsilon_{-t} \tag{4.2.6a}$$

$$y_{-t} = Gx_{-t-1} + D\epsilon_{-t} \tag{4.2.6b}$$

where ϵ_{-t} is an i.i.d. Gaussian vector with mean zero and covariance matrix I. We want a recursive way to compute the projections $\hat{x}_{-t} = E\left[x_{-t}|y^{-t}\right], \hat{y}_{-t} = E\left[y_{-t}|y^{-t-1}\right]$ where $y^{-t} \equiv [y_{-t}, y_{-t-1}, \ldots]$.⁵ Let Σ be the covariance matrix of the state-reconstruction errors $e_{-t} = x_{-t} - \hat{x}_{-t}$, conditional on y^{-t} . The maximum-likelihood estimator \hat{x}_{-t} maximizes $-e'_{-t}\Sigma^{-1}e_{-t}$. The Kalman filter constructs Σ and gives a recursive way of computing \hat{x}_{-t} as a function of the infinite history y^{-t} . In particular, the Kalman filter attains the representation

$$\hat{x}_{-t} = A\hat{x}_{-t-1} + K\left(y_{-t} - \hat{y}_{-t}\right) \tag{4.2.7a}$$

$$\hat{y}_{-t} = G\hat{x}_{-t-1} \tag{4.2.7b}$$

where K is the Kalman gain. Equations (4.2.6), (4.2.7) imply that the prediction errors satisfy $y_{-t} - \hat{y}_{-t} = G(x_{-t-1} - \hat{x}_{-t-1}) + D\epsilon_{-t}$. Define the error in estimating x_{-t} as $e_{-t} = x_{-t} - \hat{x}_{-t}$. Substitute (4.2.7) into (4.2.6) to deduce

$$e_{-t} = (A - KG) e_{-t-1} + (C - KD) \epsilon_{-t}.$$
(4.2.8)

Define the error covariance matrix $\Sigma_{-t} = Ee_{-t}e'_{-t}$. Then for a fixed, not necessarily optimal K, (4.2.8) implies

$$\Sigma_{-t} = (A - KG) \Sigma_{-t-1} (A - KG)' + (C - KD) (C - KD)'.$$
(4.2.9)

⁴ The text of this section assumes an infinite history y^t . Alternatively, let *s* denote a finite horizon. Then for the filtering problem with the timing convention of this section, we would have an initial condition stating that e_{-s} has a Gaussian distribution with mean zero and covariance matrix Σ_0 . This corresponds to setting a terminal value function for the dual control problem with the quadratic form $\lambda'_s \Sigma_0 \lambda_s$. Under the different convention about time indexes that we shall use in section 4.3 and the rest of this chapter, for the horizon *s* version of the problem, the initial condition for the filtering problem is stated in terms of a quadratic form $e'_0 \Sigma_0^{-1} e_0$. That corresponds to a terminal condition stated in terms of $\lambda'_0 \Sigma_0 \lambda_0$. It is a terminal condition because the flow of time is reversed.

⁵ Note the different conditioning information denoted by \hat{x}_{-t} and \hat{y}_{-1} .

The limit of iterations on (4.2.9) satisfies

$$\Sigma = (A - KG) \Sigma (A - KG)' + (C - KD) (C - KD)'.$$
(4.2.10)

The value of K that minimizes Σ in (4.2.10) satisfies

$$K = (CD' + A\Sigma G') (DD' + G\Sigma G')^{-1}.$$
(4.2.11)

Formulas (4.2.10), (4.2.11) implement the steady state Kalman filter. An efficient algorithm for computing (K, Σ) iterates on (4.2.11), (4.2.10), starting from the initial value $\Sigma = 0$. This is a version of the Howard policy improvement algorithm.

Equations (4.2.11), (4.2.10) also implement the policy improvement algorithm for solving a particular optimal linear regulator that is defined in terms of a state vector λ_t and a control vector μ_t . Given the initial value of the state, λ_0 , the dual problem is

$$\max_{\{\mu_t\}} \left\{ -.5 \sum_{t=0}^{\infty} \tilde{z}_t' \tilde{z}_t \right\}$$

$$(4.2.12)$$

where the maximization is subject to λ_0 given and

$$\tilde{z}_t = C'\lambda_t + D'\mu_t \tag{4.2.13a}$$

$$\lambda_{t+1} = A'\lambda_t + G'\mu_t. \tag{4.2.13b}$$

Equation (4.2.13a) defines the objective function. The solution of the optimal linear regulator is a policy rule

$$\mu_t = -K'\lambda_t \tag{4.2.14}$$

that attains the optimal value function

$$v\left(\lambda_0\right) = -.5\lambda_0'\Sigma\lambda_0. \tag{4.2.15}$$

We shall show that $\lambda_0 = \Sigma^{-1} e_0$ and that therefore the optimized value $-.5\lambda'_0 \Sigma \lambda_0$ in (4.2.12) equals the quadratic term $-.5e'_0 \Sigma^{-1} e_0$ in a log-likelihood function.

The key practical insight of these findings is that we can compute the pair (Σ, K) for the filtering problem by solving the associated optimal linear regulator (4.2.12), (4.2.14). The reversal in time and the transposition of matrices

as we move from the filtering problem to the optimal linear regulator problem are manifestations of duality, as subsequent sections show.

The duality of optimal filtering and control brings substantial insights and computational advantages. In chapters 13 and 14, we shall use these insights again to pose and solve robust filtering problems.

The remainder of this chapter substantiates our claims about duality. The reader who is willing to accept the preceding assertions about duality on faith can proceed immediately to subsequent chapters. Though it can be skipped, we think that the subsequent arguments convey some of the magic associated with the duality of filtering and control.

4.3. Sequence version of primal and dual problems

This section substantiates various assertions in the previous section. We show how the Kalman filtering problem leads to an augmented optimal linear regulator problem in terms of dual variables. We now let the time index t flow forward. This has the consequence that a reversal of time will occur in the dual problem. We consider the state space system for $t \ge 1$:

$$x_t = Ax_{t-1} + C\epsilon_t \tag{4.3.1a}$$

$$y_t = Gx_{t-1} + D\epsilon_t. \tag{4.3.1b}$$

Here ϵ_t , $t \ge 1$, is an i.i.d. Gaussian disturbance vector with mean zero and covariance matrix I. We take the initial condition x_0 to be unknown with prior distribution described by

$$x_0 = \hat{x}_0 + e_0 \tag{4.3.2}$$

where e_0 is a Gaussian vector with mean zero and covariance matrix $Ee_0e'_0 = \Sigma_0$. We assume that e_0 is distributed independently of the ϵ_t 's for $t \ge 0$. For any variable z, let z^s be the vector of observations on $\{z_t, t = 1, \ldots, s\}$. The joint density of (y^s, x^s) is Gaussian. Therefore it can be represented

$$f(x^s, y^s) \propto \exp\left(-\mathcal{D}_s\right)$$

where

$$\mathcal{D}_s = \frac{1}{2} e_0' \Sigma_0^{-1} e_0 + \frac{1}{2} \sum_{t=1}^s \epsilon_t' \epsilon_t.$$
(4.3.3)

Whittle (1990, 1996) calls D_s the 'discrepancy'. To see that the time t contribution to D_s is $(1/2)\epsilon'_t\epsilon_t$, note that by (4.3.1)

$$\begin{bmatrix} x_t \\ y_t \end{bmatrix} = \begin{bmatrix} A \\ G \end{bmatrix} x_{t-1} + C^* \epsilon_t,$$

where $C^* = \begin{bmatrix} C \\ D \end{bmatrix}$. The covariance matrix of $C^* \epsilon_t$ is $C^* C^{*'}$. Then the time t contribution to the discrepancy is⁶

$$\frac{1}{2}\epsilon_t'C^{*\prime}\left(C^*C^{*\prime}\right)^{-1}C^*\epsilon_t = \frac{1}{2}\epsilon_t'\epsilon_t.$$

4.3.1. Sequence version of Kalman filtering problem

Given y^s , we seek estimators of the hidden state x_t for $t = 1, \ldots, s - 1$. We observe y^s and estimate the hidden states by maximizing the log likelihood $-D_s$ with respect to the unobserved states and shocks ϵ_s . In particular, we seek values of $e_0, \{\epsilon_t, x_{t-1}\}_{t=1}^s$ that minimize (4.3.3) subject to (4.3.1), (4.3.2). Following Whittle (1990, 1996), we formulate this minimization problem in terms of a Lagrangian. Letting $\{\lambda_t, \mu_{t+1}\}_{t=0}^s$ be sequences of vectors of Lagrange multipliers, we form

$$J_{1} = \frac{1}{2} e_{0}' \Sigma_{0}^{-1} e_{0} + \frac{1}{2} \sum_{t=1}^{s} \epsilon_{t}' \epsilon_{t} + \lambda_{0}' (x_{0} - \hat{x}_{0} - e_{0}) + \sum_{t=1}^{s} \lambda_{t}' (x_{t} - Ax_{t-1} - C\epsilon_{t}) + \sum_{t=1}^{s} \mu_{t}' (y_{t} - Gx_{t-1} - D\epsilon_{t}).$$

$$(4.3.4)$$

⁶ The matrix $(C^*C^{*'})^{-1}C^*$ is the Moore-Penrose generalized inverse of $C^{*'}$.

4.3.2. Sequence version of dual problem

We want to minimize J_1 with respect to e_0 , ϵ_t for $t = 1, \ldots, s$, and x_t for $t = 0, \ldots, s-1$ and to maximize with respect to $\lambda_t, t = 0, \ldots, s$, and $\mu_t, t = 1, \ldots, s$. To illuminate how the Kalman filter is the dual of a linear regulator, we optimize in a particular order, thereby eventually arriving at a reduced Lagrangian that takes the form of an augmented linear regulator problem.

4.3.2.1. Minimizing over e_0, ϵ_t

Following Whittle (1990, 1996), we first minimize with respect to $e_0, \epsilon_t, t = 1, \ldots, s$. The first order conditions with respect to ϵ_t and e_0 can be written

$$\epsilon_t = C'\lambda_t + D'\mu_t \tag{4.3.5a}$$

$$e_0 = \Sigma_0 \lambda_0. \tag{4.3.5b}$$

Condition (4.3.5a) implies that

$$\epsilon'_t \epsilon_t = \begin{bmatrix} \lambda_t \\ \mu_t \end{bmatrix}' \begin{bmatrix} CC' & CD' \\ DC' & DD' \end{bmatrix} \begin{bmatrix} \lambda_t \\ \mu_t \end{bmatrix}.$$
(4.3.6)

A quick calculation also shows that

$$\lambda_t' C \epsilon_t + \mu_t' D \epsilon_t = \begin{bmatrix} \lambda_t \\ \mu_t \end{bmatrix}' \begin{bmatrix} CC' & CD' \\ DC' & DD' \end{bmatrix} \begin{bmatrix} \lambda_t \\ \mu_t \end{bmatrix}.$$
(4.3.7)

Condition (4.3.5b) implies that

$$e_0' \Sigma_0^{-1} e_0 = \lambda_0' \Sigma_0 \lambda_0 \tag{4.3.8}$$

and that

$$\lambda_0' \left(x_0 - \hat{x}_0 - e_0 \right) = \lambda_0' \left(x_0 - \hat{x}_0 - \Sigma_0 \lambda_0 \right).$$
(4.3.9)

Note the presence of Σ_0 rather than Σ_0^{-1} on the right side of (4.3.8). Substituting (4.3.6), (4.3.7), (4.3.8), and (4.3.9) into (4.3.4) gives $J_1 = J_2$ where

$$J_{2} = -\frac{1}{2}\lambda_{0}'\Sigma_{0}\lambda_{0} - \frac{1}{2}\sum_{t=1}^{s} \begin{bmatrix} \lambda_{t} \\ \mu_{t} \end{bmatrix}' \begin{bmatrix} CC' & CD' \\ DC' & DD' \end{bmatrix} \begin{bmatrix} \lambda_{t} \\ \mu_{t} \end{bmatrix} + \lambda_{0}' (x_{0} - \hat{x}_{0}) + \sum_{t=1}^{s} \lambda_{t}' (x_{t} - Ax_{t-1}) + \sum_{t=1}^{s} \mu_{t}' (y_{t} - Gx_{t-1}).$$

$$(4.3.10)$$

By expressing the objective in terms of the dual variables (i.e., the multipliers μ_t, λ_t), through equation (4.3.8) the objective function in (4.3.10) involves a quadratic form in Σ_0 rather than Σ_0^{-1} . This feature is important for understanding the duality of filtering and control.

4.3.2.2. Extremizing over $\lambda_t, \mu_t; x_t$

We want to maximize J_2 with respect to $\lambda_t, t = 0, \ldots, s$ and $\mu_t, t = 1, \ldots, s$, and to minimize it with respect to $x_t, t = 0, \ldots, s - 1$. Minimizing (4.3.10) with respect to $x_t, t = 0, \ldots, s - 1$ yields the first-order condition

$$\lambda_{t-1} = A'\lambda_t + G'\mu_t. \tag{4.3.11}$$

Having minimized out the x_t 's, we are left with the problem of choosing $\lambda_t, t = 0, \ldots, s$ and $\mu_t, t = 1, \ldots, s$ to maximize

$$J_{3} = -\frac{1}{2}\lambda_{0}'\Sigma_{0}\lambda_{0} - \frac{1}{2}\sum_{t=1}^{s} \begin{bmatrix} \lambda_{t} \\ \mu_{t} \end{bmatrix}' \begin{bmatrix} CC' & CD' \\ DC' & DD' \end{bmatrix} \begin{bmatrix} \lambda_{t} \\ \mu_{t} \end{bmatrix}$$

$$-\lambda_{0}'\hat{x}_{0} + \sum_{t=1}^{s} \mu_{t}'y_{t}$$

$$(4.3.12)$$

subject to (4.3.11) and the boundary conditions $\lambda_t = 0$, $\mu_t = 0$ for t > s. Here $J_3 = J_2$. Notice how this resembles a finite horizon *augmented linear regulator* problem (see page 60) with state vector λ_t and control vector μ_t . However, the direction of time is reversed. The term $-\frac{1}{2}(\lambda'_0\Sigma_0\lambda_0 + 2\lambda'_0\hat{x}_0)$ plays the role of a 'terminal' value function once time is reversed. The optimal control takes the form of a feedback rule

$$\mu_t = -K'_t \lambda_t + g_t y_t + f_t \hat{x}_0, \qquad (4.3.13)$$

where K_t is a version of the Kalman gain, as we shall see in detail below.

4.4. Digression: reversing the direction of time

We briefly return to a formulation of the filtering problem in which time recedes into the past with increases in t, as in section 4.2. Supposing that s > 0 and letting $t = 0, \ldots, s$, the state space system is (4.2.6) where the initial condition at time -s - 1 is

$$x_{-s-1} = \hat{x}_{-s-1} + e_{-s-1}$$

where e_{-s-1} is a Gaussian random vector with mean zero and covariance matrix Σ_{-s-1} . Define the discrepancy at horizon s as

$$D_s = \frac{1}{2} e'_{-s-1} \Sigma_{-s-1}^{-1} e_{-s-1} + \frac{1}{2} \sum_{t=0}^{s} \epsilon'_{-t} \epsilon_{-t}.$$
 (4.4.1)

We could follow the steps in the previous section to derive the dual problem with these timing conventions. In the limit as $s \to +\infty$, the dual problem would assume the form of the optimal linear regulator (4.2.12), (4.2.14).

For the remainder of this chapter, we shall use the timing conventions of section 4.3. However, in chapter 13, we shall again use the timing convention of section 4.2.

4.5. Recursive version of dual problem

We are sometimes interested in versions of problem (4.3.12) that condition on infinite histories of observations, in which case there is a recursive formulation of the problem. We seek a time invariant K, which we attain by studying the problem as $s \to \infty$ and then taking the limit of K_t as $t \to \infty$. The recursive version of problem (4.3.12) is associated with the Bellman equation

$$-\frac{1}{2}\lambda'\Sigma\lambda - \lambda'\hat{x} - \frac{1}{2}\iota = \max_{\mu,\lambda^*} \left\{ -\frac{1}{2}\lambda^{*'}\Sigma^*\lambda^* - \frac{1}{2} \begin{bmatrix} \lambda \\ \mu \end{bmatrix}' \begin{bmatrix} CC' & CD' \\ DC' & DD' \end{bmatrix} \begin{bmatrix} \lambda \\ \mu \end{bmatrix} + \mu'y - \lambda^{*'}\hat{x}_0 \right\}$$
(4.5.1)

where the maximization on the right is subject to the law of motion

$$\lambda^* = A'\lambda + G'\mu \tag{4.5.2}$$

and where λ^* now denotes *last* period's value of λ and Σ^* is last period's value of Σ . The term ι is a constant that we'll explain later. This Bellman equation induces a mapping from Σ^* to Σ . The unique positive semi-definite matrix fixed point Σ and the matrix K associated with the optimal feedback rule supply the ingredients (Σ, K) that solve the infinite-history Kalman filtering problem.

Letting ψ be a vector of Lagrange multipliers on (4.5.2), the first-order conditions with respect to λ^*, μ for maximizing (4.5.1) subject to (4.5.2) are

$$0 = -\Sigma^* \lambda^* - \hat{x}_0 - \psi$$

$$0 = -DC' \lambda + y + G\psi - DD'\mu$$

Eliminate ψ and rearrange to get the feedback rule

$$\mu = -K'\lambda + (G\Sigma^*G' + DD')^{-1} (y - G\hat{x}_0), \qquad (4.5.3)$$

where

$$K = (CD' + A\Sigma^*G') (DD' + G\Sigma^*G')^{-1}.$$
(4.5.4)

The matrix K is the Kalman gain. When (4.5.4) is evaluated at the stationary solution $\Sigma = \Sigma^*$ of the Riccati equation implied by the Bellman equation (4.5.1), (4.5.3) solves the infinite-history, time-invariant filtering problem. We now indicate how (4.5.1) implies a Riccati equation mapping Σ^* into Σ .

Use (4.5.2) and (4.5.3) to express λ^* as

$$\lambda^* = (A - KG)' \lambda + G' (G\Sigma^* G' + DD')^{-1} (y - G\hat{x}_0).$$
(4.5.5)

Using (4.5.3) and (4.5.5) to evaluate the quadratic forms in λ_0 on the first line of the right side of (4.5.1) shows

$$\left\{\lambda^{*'}\Sigma^*\lambda^* + \begin{bmatrix}\lambda\\\mu\end{bmatrix}'\begin{bmatrix}CC'&CD'\\DC'ⅅ'\end{bmatrix}\begin{bmatrix}\lambda\\\mu\end{bmatrix}\right\} = \lambda'\Sigma\lambda + \text{terms in } (y - G\hat{x}_0)$$

where

$$\Sigma = (A - KG) \Sigma^* (A - KG)' + (C - KD) (C - KD)'.$$
(4.5.6)

Formula (4.5.6) in conjunction with formula (4.5.4) is one form of the Riccati equation for the conditional covariance matrix Σ for the hidden state next period.

For the next step of the argument, we temporarily ignore the term in $y-G\hat{x}_0$ appearing in (4.5.3). Then, using (4.5.5) and $\mu = -K'\lambda$, we can calculate that

$$\mu' y - \lambda^{*'} \hat{x}_0 = -\lambda' \left(A \hat{x}_0 + K \left(y - G \hat{x}_0 \right) \right) \equiv -\lambda' \hat{x}$$
(4.5.7)

where

$$\hat{x} = A\hat{x}_0 + K\left(y - G\hat{x}_0\right) \tag{4.5.8}$$

is the estimator of the state for next period. Formulas (4.5.8) and (4.5.4), evaluated at the fixed point of (4.5.6) are the standard time-invariant Kalman filtering formulas.

Finally, we have to complete and collect the terms coming from $(G\Sigma^*G' + DD')^{-1}(y - G\hat{x}_0)$ in (4.5.3). Tedious algebra verifies that they contribute the term

$$\iota = (y - G\hat{x}_0)' (G\Sigma^* G' + DD')^{-1} (y - G\hat{x}_0)$$

that appears on the left side of (4.5.1). The matrix $G\Sigma^*G' + DD'$ is the covariance matrix of the innovations $y - G\hat{x}_0$.

4.6. Recursive version of Kalman filtering problem

For some of our future work, it is convenient to study a recursive version of the filtering problem using the dual variables again but to embrace a somewhat different perspective.

We return to the original problem. In a recursive spirit, we formulate a oneperiod filtering problem and seek a recursion in an optimized value function. The state-space system is

$$x = Ax_0 + C\epsilon \tag{4.6.1a}$$

$$y = Gx_0 + D\epsilon \tag{4.6.1b}$$

$$x_0 = \hat{x}_0 + e_0, \tag{4.6.1c}$$

where ϵ is a Gaussian random vector with mean zero and identity covariance matrix and e_0 is a Gaussian random vector distributed independently of ϵ with mean 0 and covariance matrix Σ_0 . The joint density of (x, y) is

$$f(x,y) \propto \exp\left(-\mathcal{D}\right)$$

where

$$\mathcal{D} = \frac{1}{2} \left(e_0' \Sigma_0^{-1} e_0 + \epsilon' \epsilon \right). \tag{4.6.2}$$

Given y, \hat{x}_0 , we want to choose (ϵ, x) to maximize the log likelihood, or equivalently, to minimize discrepancy \mathcal{D} subject to (4.6.1). We will show that the optimized value of the discrepancy (4.6.2) takes the form

$$\frac{1}{2}e_1'\Sigma_1^{-1}e_1 + \frac{1}{2}\iota \tag{4.6.3}$$

where $e_1 = x - \hat{x}_1$, $\hat{x}_1 = A\hat{x}_0 + K(y - G\hat{x}_0)$, K is the Kalman gain, Σ_1 is related to Σ_0 by a matrix Riccati difference equation, and ι , defined in our discussion of (4.5.1), is the contribution to the log-likelihood function (entropy) that cannot be influenced by the filter. Thus, we have the Bellman equation

$$\frac{1}{2}e_1'\Sigma_1^{-1}e_1 + \frac{1}{2}\iota = \min_{\epsilon,x} \left\{ .5e_0'\Sigma_0^{-1}e_0 + \epsilon'\epsilon \right\}$$
(4.6.4)

where the minimization is subject to (4.6.1). Further, the quadratic form $e'_1 \Sigma_1^{-1} e_1$ on the left equals the quadratic form $\lambda'_1 \Sigma_1 \lambda_1$ that appears on the left side of the Bellman equation for the dual problem (4.5.1).

To solve the filtering problem for an additional period, we would use Σ_1 to update the criterion (4.6.2) to be $\frac{1}{2} \left(e'_1 \Sigma_1^{-1} e_1 + \epsilon' \epsilon \right)$ and continue as before with next period's observation on y and $e_1 = x - \hat{x}_1$.

It is useful to solve the recursive version of the filtering problem using Lagrangian methods. Form the Lagrangian

$$J = \frac{1}{2} \left(e_0' \Sigma_0^{-1} e_0 + \epsilon' \epsilon \right) + \lambda_0' \left(x_0 - \hat{x}_0 - e_0 \right) + \lambda' \left(x - A x_0 - C \epsilon \right) + \mu' \left(y - G x_0 - D \epsilon \right).$$

The first-order conditions for minimizing J with respect to (ϵ, e_0) imply

$$\epsilon = C'\lambda + D'\mu \tag{4.6.5a}$$

$$e_0 = \Sigma_0 \left(A' \lambda + G' \mu \right),$$
 (4.6.5b)

where we are using the first-order condition $\lambda_0 = A'\lambda + G'\mu$ to get (4.6.5b).

The equality $e_0 = x_0 - \hat{x}_0$ and (4.6.1) imply

$$x - A\hat{x}_0 = C\epsilon + Ae_0 \tag{4.6.6a}$$

$$y - G\hat{x}_0 = D\epsilon + Ge_0. \tag{4.6.6b}$$

Substitute (4.6.5) into (4.6.6) and rearrange to get

$$\begin{bmatrix} y - G\hat{x}_0\\ x - A\hat{x}_0 \end{bmatrix} = \Lambda \begin{bmatrix} \mu\\ \lambda \end{bmatrix},$$
(4.6.7)

where

$$\Lambda = \begin{bmatrix} G\Sigma_0 G' + DD' & DC' + G\Sigma_0 A' \\ CD' + A\Sigma_0 G' & A\Sigma_0 A' + CC' \end{bmatrix}.$$
(4.6.8)

Then

$$\begin{bmatrix} \mu \\ \lambda \end{bmatrix} = \Lambda^{-1} \begin{bmatrix} y - G\hat{x}_0 \\ x - A\hat{x}_0 \end{bmatrix}.$$

For reasons to be explained in chapter 14, we call the optimized value of $\epsilon' \epsilon + e'_0 \Sigma_0^{-1} e_0$ the conditional entropy of (y, x) and denote it $\operatorname{ent}(y, x)$. It is the maximized value of the log likelihood function. Using (4.6.5), we can evaluate $\operatorname{ent}(y, x)$ to be

ent
$$(y, x) \equiv \epsilon' \epsilon + e'_0 \Sigma_0^{-1} e_0 = \begin{bmatrix} \mu \\ \lambda \end{bmatrix}' \Lambda \begin{bmatrix} \mu \\ \lambda \end{bmatrix}$$

$$= \begin{bmatrix} y - G\hat{x}_0 \\ x - A\hat{x}_0 \end{bmatrix}' \Lambda^{-1} \begin{bmatrix} y - G\hat{x}_0 \\ x - A\hat{x}_0 \end{bmatrix}.$$
 (4.6.9)

Let

$$L = \begin{bmatrix} I & 0 \\ -K & I \end{bmatrix}$$

where

$$K = \Lambda_{21}\Lambda_{11}^{-1} \equiv \left(A\Sigma_0 G' + CD'\right) \left(DD' + G\Sigma_0 G'\right)^{-1}.$$
 (4.6.10)

We recognize K to be the Kalman gain. It can be verified that

$$L\Lambda L' = \begin{bmatrix} \Lambda_{11} & 0\\ 0 & \Lambda_{22} - \Lambda_{21}\Lambda_{11}^{-1}\Lambda'_{21} \end{bmatrix},$$
(4.6.11)

where

$$\Lambda_{11} = G\Sigma_0 G' + DD' \tag{4.6.12}$$

and

$$\Sigma_{1} \equiv \Lambda_{22} - \Lambda_{21}\Lambda_{11}^{-1}\Lambda_{21}'$$

= $CC' + A\Sigma_{0}A' - (A\Sigma_{0}G' + CD')(DD' + G\Sigma_{0}G')^{-1}(A\Sigma_{0}G' + CD')'.$
(4.6.13)

It turns out that Λ_{11} is the covariance matrix of the innovations $y - G\hat{x}_0$ and $\Lambda_{22} - \Lambda_{21}\Lambda_{11}^{-1}\Lambda'_{21}$ is the covariance matrix of $x - \hat{x}_1$ where \hat{x}_1 is the estimator of the state x. In particular, notice that

$$L\begin{bmatrix} y - G\hat{x}_0\\ x - A\hat{x}_0 \end{bmatrix} = \begin{bmatrix} y - G\hat{x}_0\\ x - A\hat{x}_0 - K(y - G\hat{x}_0) \end{bmatrix} = \begin{bmatrix} y - G\hat{x}_0\\ x - \hat{x}_1 \end{bmatrix}$$

where

$$\hat{x}_1 = A\hat{x}_0 + K\left(y - G\hat{x}_0\right). \tag{4.6.14}$$

Here \hat{x}_1 is the estimate of the state next period, based on the observed value of y. Thus, returning to (4.6.9), we have

$$\operatorname{ent}(y,x) = \begin{bmatrix} y - G\hat{x}_{0} \\ x - A\hat{x}_{0} \end{bmatrix}' L' (L\Lambda L')^{-1} L \begin{bmatrix} y - G\hat{x}_{0} \\ x - A\hat{x}_{0} \end{bmatrix}$$
$$= \begin{bmatrix} y - G\hat{x}_{0} \\ x - \hat{x}_{1} \end{bmatrix}' \begin{bmatrix} \Lambda_{11} & 0 \\ 0 & \Lambda_{22} - \Lambda_{21}\Lambda_{11}^{-1}\Lambda'_{21} \end{bmatrix}^{-1} \begin{bmatrix} y - G\hat{x}_{0} \\ x - \hat{x}_{1} \end{bmatrix}$$
$$= (y - G\hat{x}_{0})' \Lambda_{11}^{-1} (y - G\hat{x}_{0}) + (x - \hat{x}_{1})' (\Lambda_{22} - \Lambda_{21}\Lambda_{11}^{-1}\Lambda'_{21})^{-1} (x - \hat{x}_{1})$$
$$= (y - G\hat{x}_{0})' \Lambda_{11}^{-1} (y - G\hat{x}_{0}) + e'_{1}\Sigma_{1}^{-1}e_{1}.$$
(4.6.15)

Formula (4.6.15) inspires the updating formula (4.6.13) for the covariance matrix of $x - \hat{x}_1$. The entropy-minimizing choice of x is evidently \hat{x}_1 ; the value of y is observed, and the value \hat{x}_0 is given, so the first term on the last line of (4.6.15) cannot be influenced by the filter. It contributes ι in (4.6.3).

4.7. Concluding remarks

In the filtering and control problems of this chapter, the decision maker assumes that his state-space model is correctly specified. Later chapters extend the duality between filtering and control to filtering problems in which the decision maker fears that the model (4.2.6) is misspecified. Chapters 6 and 7 formulate and solve a robust control problem. Chapter 13 then exploits duality to discover a corresponding robust filtering problem. Effectively, that chapter works backwards from a robust version of the optimal linear regulator problem (4.2.12), (4.2.14) to get a corresponding filtering problem. Not surprisingly in view of the time-reversal between the dual and original problems, the objective function of the decision maker in the dual problem is backward-looking. While interesting, that is not always the most natural formulation for economic problems. Therefore, in chapter 14 we alter the objective function of the decision maker to be forward-looking. That leads us to another robust filtering problem. We cast that forward-looking robust filtering problem as a model-approximation problem using entropy to measure model misspecification. This forward-looking robust filtering problem has the same mathematical structure as the one studied in section 4.6. In chapters 6 and 7, we shall also use duality theory extensively to formulate our basic robust dynamic decision problem.

Part III Robust control and applications

Chapter 5 Static multiplier and constraint games

5.1. Introduction

By considering static examples, this chapter takes a detour from the main concern in this book, which is devising decision rules that are robust to misspecified dynamics. To simplify some of the analytical issues, this chapter strips off all dynamics and focuses on two types of interrelated static zero-sum two-player games whose equilibria induce robust decisions for the maximizing player within a one-period setting. We call them a multiplier game and a constraint game. We take up dynamic versions of both of these games in subsequent chapters.

We begin with a simple static Phillips curve example in section 5.2. Subsequent sections then focus on another simple example with the aim of exposing the role of technical assumptions that reconcile outcomes from alternative games.

We consider two classes of possible misspecifications to a static Gaussian approximating model that might concern the decision maker. The more restricted setting allows misspecifications only in the mean of a Gaussian random variable. The more generous setting allows misspecifications in the form of arbitrary alternative distributions that are absolutely continuous with respect to the approximating model. For a Gaussian approximating model, the worse case model from this class remains Gaussian, but it has distortions to both the mean and the variance.¹

¹ Chapter 2 described two related such classes of distortions for dynamic models.

5.2. Phillips curve example

To illustrate basic ideas, this section adapts Kydland and Prescott's (1977) model of a policy maker who sets inflation in view of an expectational Phillips curve. We modify of Kydland and Prescott's model² by assuming that the policy maker views his model as an approximation. The policy maker solves a *multiplier* game as a way to compute a decision that is robust to model misspecification. Let U, π, π_e be the unemployment rate, the inflation rate, and the public's expected rate of inflation, respectively. The government's approximating model is

$$U = U^* - \gamma \left(\pi - \pi_e\right) + \hat{\epsilon} \tag{5.2.1}$$

where $\gamma > 0$ and $\hat{\epsilon}$ is $\mathcal{N}(0,1)$. Here U^* is the natural rate of unemployment, the unemployment rate that on average prevails when $\pi = \pi_e$.³ The government sets π , the public sets π_e , and nature draws $\hat{\epsilon}$. The government views (5.2.1) as an approximation in the sense that it suspects that U might actually be governed by

$$U = U^* - \gamma (\pi - \pi_e) + (\epsilon + w), \qquad (5.2.2)$$

where ϵ is another random variable that is distributed $\mathcal{N}(0,1)$ and w is an unknown distortion to the mean ϵ . Thus, the government suspects that the natural unemployment rate might be $U^* + w$ for some unknown w. The government does know that

$$w^2 \le \eta. \tag{5.2.3}$$

The parameter η bounds the square of the government's specification error w^2 .

 $^{^2}$ We are building on Sargent's (1999) rendition of Kydland and Prescott's model in the style of Stokey (1989XX).

³ To bring the setup closer to that used in dynamic settings in chapters 2 and 6, we could have added a parameter c and expressed (5.2.2) as $U = U^* - \gamma(\pi - \pi_e) + c(\epsilon + w)$, where c is used to scale the volatility of the noise ϵ . We have set c = 1 to simplify some formulas in this chapter.

5.2.1. The government's problem

The government values outcomes (U, π) according to the utility function assigned by Kydland and Prescott, namely,

$$-E\left(U^2 + \pi^2\right) \tag{5.2.4}$$

where E denotes the mathematical expectation. Because it does not trust the approximating model, the government cares about the mathematical expectation over multiple models indexed by w's that satisfy (5.2.3).

We proceed in the spirit of Stokey's (1989) analysis of credible government policies. We derive the government's robust best response to the private sector's setting of π_e . The appendix then uses that robust best response function to formulate a rational expectations equilibrium. The government's best response function takes π_e as fixed. Given π_e , the government wants to set π so that it attains satisfactory outcomes for all $w^2 \leq \eta$. The government therefore sets π equal to the equilibrium π -component of the following two-player zero-sum multiplier game

$$\max_{\pi} \min_{w} -E\left\{U^2 + \pi^2\right\} + \theta w^2 \tag{5.2.5}$$

where both the minimization and maximization are subject to (5.2.2) and $\theta > 1$ is a fixed penalty parameter. We shall soon explain how the penalty parameter θ relates to η in (5.2.3) and why we impose $\theta > 1$. We shall also discuss conditions that let us exchange the order of maximization and minimization in (5.2.5). The first order conditions for π and w, respectively, for problem (5.2.5) are

$$(1+\gamma^2)\pi - \gamma^2\pi_e - \gamma (U^* + w) = 0$$
(5.2.6a)

$$U^* - \gamma \pi + \gamma \pi_e + w (1 - \theta) = 0.$$
 (5.2.6b)

Solving these equations jointly for π, w as functions of π_e gives:

$$\pi\left(\theta\right) = \left(\frac{\gamma}{1 - \theta^{-1} + \gamma^2}\right) \left(U^* + \gamma \pi_e\right) \tag{5.2.7}$$

$$w\left(\theta\right) = \left(\frac{\theta^{-1}}{1 - \theta^{-1} + \gamma^2}\right) \left(U^* + \gamma \pi_e\right).$$
(5.2.8)

Here $\pi(\theta)$ gives the government's (robust) best response function for setting π as a function of π_e , while $w(\theta)$ determines the worst case model, given π_e and the government's setting $\pi(\theta)$.

Note that when $\theta = +\infty$ so that there is no concern for model misspecification,

$$\pi(\infty) = \left(\frac{\gamma}{1+\gamma^2}\right) (U^* + \gamma \pi_e) \tag{5.2.9}$$

$$w\left(\infty\right) = 0. \tag{5.2.10}$$

Note also that (5.2.6a) says that $\pi(\theta)$ satisfies

$$\pi\left(\theta\right) = \left(\frac{\gamma}{1+\gamma^{2}}\right)\left[\left(U^{*}+w\left(\theta\right)\right)+\gamma\pi_{e}\right].$$

This equation defines a function

$$\pi\left(\theta\right) = B\left(\pi_e;\theta\right),\tag{5.2.11}$$

which is the government's robust best response function to the state of expectations π^e . Evidently the robust rule can be obtained by replacing the estimate of the natural unemployment rate U^* under the approximating model in (5.2.9) with the worst case estimate of the natural rate $U^* + w(\theta)$. Thus, one way to achieve robustness is to distort estimates of exogenous variables in a pessimistic way relative to the approximating model, then to proceed with ordinary decision making procedures.⁴ A related characterization of robust decision making procedures will prevail in the dynamic settings to be studied in subsequent chapters. However, because the models there are dynamic, the distortions become more interesting and involve misspecifications in how state vectors feed back on their own histories.

It is useful to compute the limiting decision $\pi(\theta)$ and worst case distortion $w(\theta)$ as $\theta \searrow 1$:⁵

$$\pi(1) = \gamma^{-1}U^* + \pi_e \tag{5.2.12}$$

$$w(1) = \gamma^{-2} \left(U^* + \gamma \pi_e \right). \tag{5.2.13}$$

In the appendix to this chapter we show how the unit slope of the government's best response to π_e in (5.2.12) will cause a rational expectations equilibrium

 $^{^4}$ See the citation attributed to Fellner on page 32.

⁵ The value $\theta = 1$ is the breakdown point to be discussed later. In the generalization of the model where $c(\epsilon + w)$ replaces $(\epsilon + w)$, the breakdown point is $\theta = c^2$.

inflation rate to approach $+\infty$ as $\theta \searrow 1$. That rational expectations inflation rate satisfies $\pi = \pi_e$, as well as having π be a robust best government response to π_e .

Given π_e , we can now tell how the penalty parameter θ is related to the constraint η . The multiplier game is in effect the Lagrangian associated with a closely related *constraint game*:

$$\sup_{\pi} \inf_{|w| \le \sqrt{\eta}} -E\left(U^2 + \pi^2\right),$$

where θ will turn out to be the Lagrange multiplier on the constraint $w^2 \leq \eta$. The associated Lagrangian is

$$\sup_{\pi} \sup_{\theta \ge 0} \inf_{w} -E\left(U^2 + \pi^2\right) + \theta\left(w^2 - \eta\right)$$

If $\theta > 1$, for the inner minimization part of this problem it is evidently optimal to set w so that the constraint $w^2 \leq \eta$ holds with equality. Then set $w = \pm \sqrt{\eta}$ and solve (5.2.8) for θ :

$$\theta = 1 + \frac{|U^* - \gamma (\pi - \pi_e)|}{\sqrt{\eta}}.$$
 (5.2.14)

Equation (5.2.14) shows how to map η into an associated θ . As described by equation (5.2.14), the parameter θ thus measures the set of alternative models over which the decision maker seeks a satisfactory outcome. We shall discuss the connection between the constraint game and the multiplier game further in the following sections. Before that, we briefly describe the sense in which (5.2.7) gives a decision for π that is robust to model misspecification.

5.2.2. Robustness of robust decisions

For convenience, we define $\sigma = -\theta^{-1}$; σ is the risk-sensitivity parameter of Jacobson (1973) and Whittle (1990). Fig. 5.2.1 illustrates the sense in which a robust decision for π is robust. Let $J(\sigma_1, \sigma_2)$ be the value of $-E(U^2 + \pi^2)$ associated with setting $\pi = \pi(\sigma_1)$ when $w = w(\sigma_2)$. Assuming $\gamma = 1, U^* = 5$, for three settings of inflation $\pi(\sigma_1)$, Fig. 5.2.1 plots $J(\sigma_1, \cdot)$ as a function of σ_2 , where the worst case $w = w(\sigma_2)$ varies along the ordinate axis. Notice how the three payoff functions $\mathcal{J}(\sigma_1, \cdot)$ cross. The $\sigma = \sigma_1 = 0$ rule gives the highest value for the government's objective when there is no specification error (i.e., $\sigma_2 = 0$ implies that w = 0), but its performance deteriorates more quickly than the robust ($\sigma_1 = -.25, \sigma_1 = -.5$) rules as w increases along the σ_2 axis. The robust rules sacrifice performance when the approximating model is correct. However, they experience lower rates of deterioration in the objective \mathcal{J} as the specification error increases.



Figure 5.2.1: Values of $\mathcal{J}(\sigma_1, \sigma_2) = -E(U^2 + \pi^2)$ for three decision rules $\pi(\sigma_1)$ for $\sigma_1 = 0, -.25, -.5$ for the worst-case $w(\sigma_2)$ for values of σ_2 on the ordinate axis. The $\sigma_1 = 0$ rule works best when w = 0, but its performance deteriorates more rapidly as |w| increases than do the robust rules.

Because our principal focus in this chapter is single-agent robust control theory, we have taken π_e as given. To complete the analysis of the Kydland-Prescott model, we should describe how π_e is set. Appendix A applies the notion of a rational expectations equilibrium to make π_e equal to the $\pi(\sigma)$ chosen by the robust monetary authority. We postpone that material to the appendix because it involves issues that would interrupt our main line of argument. We now turn to important technical details about our single agent decision model.

5.3. Basic setup with a correct model

This section uses a very simple static model to describe in more detail the relationship between a static constraint game and a static multiplier game. Let x be an endogenous variable and u a scalar control variable. The variables u and x are linked by the approximating model

$$x = u + \hat{\epsilon} \tag{5.3.1}$$

where $\hat{\epsilon}$ is a random variable with mean zero and variance 1. Letting *E* denote the mathematical expectation and *b* be a scalar, a decision maker wants (u, x) to maximize

$$-\frac{u^2}{2} - \frac{1}{2}E(x-b)^2$$
(5.3.2)

or

$$-\frac{u^2}{2} - \frac{(u-b)^2}{2} - \frac{1}{2}.$$
 (5.3.3)

The maximizing choice is $u = \frac{b}{2}$.

We want to think about the situation where the decision maker treats the model (5.3.1) not as true but as an approximation. To represent specification error, the decision maker replaces the approximating model (5.3.1) with the distorted model

$$x = u + (\epsilon + w), \qquad (5.3.4)$$

where ϵ is another random variable with mean zero and variance 1. The distorted model thus has a random term with unknown mean w and known variance 1, rather than known mean 0 and variance 1 as under the approximating model (5.3.1). The decision maker formulates the idea that his model is a good approximation by assuming that $|w| \leq \sqrt{\eta}$ where $\eta > 0$. Substituting (5.3.4) into (5.3.2), the criterion function becomes

$$-\frac{u^2}{2} - \frac{\left(u+w-b\right)^2}{2} - \frac{1}{2}.$$
 (5.3.5)

The decision maker seeks a u that works well for any $|w| \leq \sqrt{\eta}$. Since the variance 1 is constant, we can replace (5.3.5) with

$$-\frac{u^2}{2} - \frac{\left(u+w-b\right)^2}{2}.$$
 (5.3.6)

Within this simple setting, we consider two types of zero-sum two-person games that can be used to choose a u that is robust to misspecifications that take the form of alternative values of w. The two games are: (1) a 'constraint game' that constrains the choices of u, v in (5.3.6) by $|w| \leq \sqrt{\eta}$; and (2) a 'multiplier game' that appends to the right side of (5.3.6) a penalty term $\frac{\theta}{2}(w^2 - \eta)$. For an appropriate choice of θ , these two formulations are equivalent under conditions identified by the Lagrange multiplier theorem (see Luenberger (1969), pp. 216-221). However, that equivalence breaks down when $\sqrt{\eta} > |b|$. As a vehicle for exploring conditions for the equivalence between the two approaches, we start by analyzing the pathological b = 0 case. Later parts of this chapter shed further light on the pathological case by allowing a larger class of misspecifications.

5.4. The constraint game with b = 0

This section considers a pathological case in which variations in the decision maker's concern about robustness, as measured by the penalty parameter θ , have no effect on his decision u. To generate the pathology, we temporarily set b = 0. To induce a robust decision u we formulate a 'constraint game':⁶

$$\max_{u} \min_{\|w\| \le \sqrt{\eta}} -\frac{u^2}{2} - \frac{(u+w)^2}{2}.$$
(5.4.1)

 $^{^{6}}$ We thank Dirk Bergemann for suggesting this example and its consequences.

Notice that the objective is concave and not convex in w (this is also true when $b \neq 0$). Also notice the timing protocol implicit in the order of maximization and minimization in (5.4.1): the maximizing player chooses first, the minimizing player second.

The equilibrium of this zero-sum two-person game can be computed by considering three possible sets of values for u. If u = 0, $w = \pm \sqrt{\eta}$ solves the inner minimization problem, with a minimized value of $-\frac{\eta}{2}$. If u > 0, the solution of the inner problem is to set $w = \sqrt{\eta}$, which makes the objective smaller than $-\frac{\eta}{2}$. Similarly, if $u \leq 0$, the solution of the inner problem is to set $w = -\sqrt{\eta}$, and the objective (5.4.1) is again smaller than $-\frac{\eta}{2}$. Thus the 'robust' decision is to set u to zero; this decision is supported by the maximizing player's expectation that w will respond to u by the rule $w = \frac{u}{|u|}\sqrt{\eta}$ for $u \neq 0$ and $w = \pm\sqrt{\eta}$ when u is zero.

A strange feature of (5.4.1) is that a preference for robustness to model misspecification has no effect on the decision u. The equilibrium outcome for u is 0, independently of the value of η .

For various reasons to be explained below, we would like to be able to exchange the order of minimization and maximization in (5.4.1). However, another peculiarity of (5.4.1) is that we cannot exchange orders of the minimization and maximization operations; neither $u = 0, w = \sqrt{\eta}$ nor $u = 0, w = -\sqrt{\eta}$ is a Nash equilibrium of the game with the order of maximization and minimization exchanged. In fact, there is no pure strategy Nash equilibrium. We will compute mixed strategy equilibria later.

5.5. Multiplier game with b = 0

We want to understand the connection between the constraint game (5.4.1) and an associated 'multiplier game'. To do so, in this section we study a Lagrangian formulation of the constraint game. This will eventually lead us to a multiplier game. We reformulate the constraint in (5.4.1) as $w^2 \leq \eta$ and form a Lagrangian:

$$\max_{u} \inf_{w} \sup_{\theta \ge 0} -\frac{u^{2}}{2} - \frac{(u+w)^{2}}{2} + \frac{\theta}{2} \left(w^{2} - \eta^{2} \right)$$

or

128

$$\max_{u} \sup_{\theta \ge 0} \inf_{w} -\frac{u^2}{2} - \frac{(u+w)^2}{2} + \frac{\theta}{2} \left(w^2 - \eta^2\right).$$
 (5.5.1)

The standard sufficient conditions for the Lagrange Multiplier Theorem do not hold here. While the constraint set for w is convex, the objective is also convex in w. As we will see in chapter 6, with appropriate qualifications, a modified version of the Lagrange Multiplier Theorem does apply.

Consider the inner minimization problem of (5.5.1), holding fixed θ and u. Suppose $\theta \leq 1$. Then the objective is convex in w (it is affine for $\theta = 1$), and the infimum over w is $-\infty$. Therefore, we need consider only $\theta > 1$. For $\theta > 1$, the first-order conditions for w are:

$$(\theta - 1)w - u = 0,$$

or

$$w = \frac{u}{\theta - 1}.$$

Consider next the second inner-most maximization problem in (5.5.1). Provided that $u \neq 0$, the supremum over θ is attained by setting θ so that the constraint is satisfied. Thus

$$\theta = 1 + \frac{|u|}{\sqrt{\eta}}$$

and

$$w = \frac{u}{|u|}\sqrt{\eta}.$$

At these values of θ, v , the objective for the outer maximization problem in (5.5.1) becomes

$$L(u) = -\frac{u^2}{2} - \frac{u^2}{2} \left(\frac{|u| + \sqrt{\eta}}{|u|}\right)^2 = -\frac{u^2}{2} - \frac{\left(|u| + \sqrt{\eta}\right)^2}{2} < -\frac{\eta}{2}$$

By making u arbitrarily close to zero, we find that the right side of the above inequality is a least upper bound. In fact, if u = 0, then w = 0 and

$$\sup_{\theta > 1} \inf_{w} -\frac{u^{2}}{2} - \frac{\left(u+w\right)^{2}}{2} + \frac{\theta}{2}\left(w^{2} - \eta\right) = -\frac{1}{2}\eta.$$

This gives the correct value of the objective of the constraint game (5.4.1), and u = 0 is the correct robust action for that game. Since the solution is not

attained at $\theta = 1$, the solution of (5.5.1) must be computed as a limit of the solution as $\theta \searrow 1$. The value $\theta = 1$ corresponds to what we shall refer to in chapter 6 as a 'breakdown point' for θ .

Games (5.4.1) and (5.5.1) are pathological because neither η in the constraint game nor θ in the multiplier game affects the equilibrium decision u. We show below how this pathology occurs because $|b| < \sqrt{\eta}$.

5.6. The model with $b \neq 0$

By setting $b \neq 0$, we can repair the pathological outcome that variations in the multiplier θ in (5.5.1) do not change the action u. We alter (5.4.1) to be:

$$\max_{u} \min_{|w| \le \sqrt{\eta}} -\frac{u^2}{2} - \frac{\left(u+w-b\right)^2}{2}.$$
(5.6.1)

The Lagrangian is:

$$\max_{u} \inf_{w} \sup_{\theta > 0} -\frac{u^{2}}{2} - \frac{\left(u + w - b\right)^{2}}{2} + \frac{\theta\left(w^{2} - \eta\right)}{2}$$

or

$$\max_{u} \sup_{\theta \ge 0} \inf_{w} -\frac{u^{2}}{2} - \frac{(u+w-b)^{2}}{2} + \frac{\theta(w^{2}-\eta)}{2}.$$

It is again true that for $\theta \leq 1$, the inner-most minimization problem has a criterion equal to $-\infty$ for any u. Thus $\theta = 1$ remains a 'breakdown point'. Variations of θ for $\theta > 1$ will now affect the decision u, thereby capturing how a concern for robustness affects the decision.

For $\theta > 1$, consider the multiplier game:

$$\max_{u} \min_{w} -\frac{u^2}{2} - \frac{\left(u+w-b\right)^2}{2} + \frac{\theta\left(w^2-\eta\right)}{2}.$$
 (5.6.2)

The objective is concave in u and convex in w. It can be verified directly that the order of maximization and minimization does not matter, and that the Nash equilibrium of the game defined by (5.6.2) can be obtained by stacking

and solving first-order conditions for the minimizing and maximizing players.⁷ The first-order conditions are:

$$u + (u + w - b) = 0$$
$$(u + w - b) - \theta w = 0.$$

The equilibrium outcomes are:

$$u = \frac{\theta b}{2\theta - 1}$$

$$w = \frac{-b}{2\theta - 1}.$$
(5.6.3)

We have thus established:

Theorem 5.6.1. For $\sqrt{\eta}$ in the interval (0, |b|) we can find a value of $\theta > 1$ for which the solution to the multiplier game (5.6.2) is the same as that of the constraint game (5.6.1) and conversely. This mapping breaks down when $\theta = 1$ and $\sqrt{\eta} \ge |b|$.

Proof. From (5.6.3), as θ ranges from $+\infty$ to one, the solution for w ranges from zero to -b.

Notice that $u = \frac{b}{2}$ for the limiting $\theta = +\infty$ case, and that u converges to b as θ declines to one.

5.6.1. Analysis of pathology

Consider now the constraint game when $\sqrt{\eta} > |b|$. Form two quadratic functions:

$$p_{-}(u) = -\frac{u^{2}}{2} - \frac{\left(u - \sqrt{\eta} - b\right)^{2}}{2}$$
$$p_{+}(u) = -\frac{u^{2}}{2} - \frac{\left(u + \sqrt{\eta} - b\right)^{2}}{2}.$$

The robust choice of u solves:

$$\max_{u} \min\{p_{-}(u), p_{+}(u)\}.$$

Notice that $p_{-}(b) = p_{+}(b)$. Moreover, $dp_{-}(0)/du = b + \sqrt{\eta}$ and $dp_{+}(0)/du = b - \sqrt{\eta}$. Because $\sqrt{\eta} > |b|$, these derivatives have opposite signs, implying that u = b remains the robust solution for large enough values of $\sqrt{\eta}$.

⁷ This is a version of von Neumann's Minimax Theorem. For example, see Dantzig (1998, pp. 286–287).



Figure 5.6.1: The functions $p_{-}(u)$, $p_{+}(u)$, $\min\{p_{-}(u), p_{+}(u)\}$ for $\sqrt{\eta} = .3$, b = 0. The maximum of $\min\{p_{-}(u), p_{+}(u)\}$ occurs at u = b = 0, a kink point of the function.



Figure 5.6.2: The functions $p_{-}(u), p_{+}(u), \min\{p_{-}(u), p_{+}(u)\}$ for $\sqrt{\eta} = .3, b = .5$. The maximum of $\min\{p_{-}(u), p_{+}(u)\}$ occurs at $u = \frac{b+\sqrt{\eta}}{2} = .4$, where the function is differentiable.

Figures 5.6.1 and 5.6.2 help reveal what is going on in the two cases $\sqrt{\eta} > |b|$ and $\sqrt{\eta} < |b|$. Fig. 5.6.1 plots the function $\min\{p_{-}(u), p_{+}(u)\}$ for $\sqrt{\eta} = .3, b =$ 0 while Fig. 5.6.2 plots it for $\sqrt{\eta} = .3, b = .5$. In Fig. 5.6.1, which corresponds to a pathological case in which $\sqrt{\eta} > |b|$, $\min\{p_{-}(u), p_{+}(u)\}$ has a maximum at u = b = 0, a nondifferentiable point formed by the intersection of the $p_{-}(u)$ and $p_{+}(u)$. In Fig. 5.6.2, for which $\sqrt{\eta} < |b|$, the maximum of $\min\{p_{-}(u), p_{+}(u)\}$ occurs at $u = \frac{\sqrt{\eta}+b}{2} = .4$, a point where the function is differentiable. Here u depends on η , reflecting a concern for robustness that was absent in the pathological $\sqrt{\eta} > |b|$ case.

5.7. Probabilistic formulation (b=0)

We now alter game (5.4.1) by enlarging the class of allowable perturbations to include more than just mean shifts. In particular, we now allow random perturbations to the approximating model. The approximating model is now

$$x = u + \epsilon$$

where $\epsilon \sim f_o(\epsilon)$ and f_o is the standard normal density. The distorted models have $\epsilon \sim f(\epsilon)$ for some density $f \neq f_o$. Corresponding to the b = 0 case above, we now let the objective in our zero-sum two-player games be

$$-\frac{u^2}{2} - \frac{\int \left(u+\epsilon\right)^2 f\left(\epsilon\right) d\epsilon}{2}.$$
(5.7.1)

To measure model misspecification we use relative entropy, which is defined to be the expected log likelihood ratio, where the expectation is evaluated at the distorted model:

$$\mathcal{I}(f) = \int \left[\log f(\epsilon) - \log f_o(\epsilon)\right] f(\epsilon) \, d\epsilon.$$
(5.7.2)

This entropy measure is convex in f. We study the game:

$$\max_{u} \min_{f,\mathcal{I}(f) \le \xi, \int f = 1} -\frac{u^2}{2} - \frac{\int (u+\epsilon)^2 f(\epsilon) d\epsilon}{2}.$$
(5.7.3)

The objective in (5.7.3) is linear in the density f and the constraint set is convex. Therefore, Lagrangian methods apply.
5.7.1. Gaussian perturbations

Before relating game (5.7.3) to game (5.4.1), we calculate the entropy measure (5.7.2) where f is a normal density with mean w and variance σ^2 . Then⁸

$$\mathcal{I}(f) = \frac{w^2}{2} + \frac{\sigma^2 - 1}{2} - \frac{\log \sigma^2}{2}.$$
(5.7.4)

Thus entropy decomposes into a part $\frac{w^2}{2}$ due to a mean distortion and a part $\frac{\sigma^2-1}{2} - \frac{\log \sigma^2}{2}$ due to a variance distortion. Because the logarithm is a concave function, the variance distortion is nonnegative:

$$\frac{\sigma^2 - 1}{2} - \frac{\log \sigma^2}{2} \ge 0.$$

To understand how game (5.7.3) relates to game (5.4.1), consider a perturbed density f that is normal with mean w and unit variance $\sigma^2 = 1$ so that the distortion consists solely of a mean shift. Then $\mathcal{I}(f) = w^2/2$ and the objective (5.7.1) becomes

$$-\frac{u^2}{2} - \frac{(u+w)^2 + 1}{2},$$

which matches (5.3.5) when b = 0. With the Gaussian $f(\epsilon)$, we can view (5.7.3) as extending (5.4.1) to a larger set of perturbations. In effect, (5.4.1) admits only perturbations that are equivalent to mean shifts in a standard normal distribution. The η in (5.4.1) relates to the parameter ξ in (5.7.3) through the formula:

$$\eta = \frac{\xi}{2}.\tag{5.7.5}$$

In shifting the distortions from numbers w to densities f, we have made the objective function linear in the distortion. The family of normal distributions with a unit variance and mean w is not convex, however. An approach that we might have but did not take is to mix w actions by allowing finite mixtures of normal distributions. Rather than doing that, we entertain more than just finite normal mixtures by allowing arbitrary densities; but we constrain their relative entropy, which effectively restricts those densities to be absolutely continuous with respect to the approximating model.

⁸ Simple calculations show that $\mathcal{I}(f)$ is the expectation of $\log(\sigma^{-1}) - (2\sigma^2)^{-1}(\epsilon - w)^2 + (2)^{-1}\epsilon^2$ evaluated with respect to $f(\epsilon)$.

5.7.2. Letting the minimizing agent make random perturbations when b = 0

By appropriately choosing f, which is now the counterpart to w in (5.4.1), the minimizing player can in effect implement a mixed strategy. This changes the solution to the problem in a substantial way.

The Lagrange saddle-point problem is:

$$\max_{u} \min_{f, \int f=1} \sup_{\theta \ge 0} -\frac{u^2}{2} - \frac{\int (u+\epsilon)^2 f(\epsilon) d\epsilon}{2} + \theta \left[\mathcal{I}(f) - \xi \right]$$

or

$$\max_{u} \max_{\theta \ge 0} \inf_{f, \int f=1} -\frac{u^2}{2} - \frac{\int (u+\epsilon)^2 f(\epsilon) d\epsilon}{2} + \theta \left[\mathcal{I}(f) - \xi \right].$$
(5.7.6)

The first-order conditions for the inner-most minimization problem of (5.7.6) are

$$\theta \left[\log f\left(\epsilon\right) - \log f_o\left(\epsilon\right) + 1 \right] + \kappa = \frac{\left(u + \epsilon\right)^2}{2}$$
(5.7.7)

where κ is a constant introduced by the constraint $\int f = 1$. The solution to this problem is:

$$f_{\theta}(\epsilon) \propto \exp\left[\frac{\left(u+\epsilon\right)^2}{2\theta}\right] f_o(\epsilon)$$
 (5.7.8)

where the constant of proportionality is chosen so that $f_{\theta}(\epsilon)$ integrates to unity. Such a constant will exist only when

$$\int \exp\left[\frac{\left(u+\epsilon\right)^2}{2\theta}\right] f_o\left(\epsilon\right) d\epsilon < \infty.$$

The integral is finite provided that $\theta > 1$. When $\theta > 1$, the density f_{θ} defined by (5.7.8) is normal since it is the product of exponentials with quadratic terms in ϵ . It is easy to verify that the density f_{θ} is proportional to the exponential of the following term:

$$\frac{\left(u+\epsilon\right)^2}{2\theta} - \frac{\epsilon^2}{2} = -\frac{\left(\theta-1\right)\epsilon^2}{2\theta} + \frac{u\epsilon}{\theta} + \frac{u^2}{2\theta}$$
$$= -\frac{\left(\epsilon-\mu_{\theta}\right)^2}{2\sigma_{\theta}^2} + c$$

where c does not depend on ϵ and where

$$\mu_{\theta} = \frac{u}{\theta - 1}$$
$$\sigma_{\theta}^2 = \frac{\theta}{\theta - 1}.$$

Thus f_{θ} is normal with mean μ_{θ} and variance σ_{θ}^2 . Notice that the variance σ_{θ}^2 becomes arbitrarily large as θ approaches unity. As a consequence, the relative entropy associated with a θ that approaches unity becomes arbitrarily large. For instance, when u = 0 (5.7.4) implies

$$\mathcal{I}(f_{\theta}) = \frac{\sigma_{\theta}^2 - 1}{2} - \frac{\log \sigma_{\theta}^2}{2}$$

There is a multiplier θ associated with each positive $\xi = 2\eta$ defined in (5.7.5). The optimized choice of u remains zero in this example, and the worst case distribution f has an increased variance (relative to the standard normal distribution) that depends on ξ . Thus, in contrast to the deterministic game, values of $\theta > 1$ correspond to specific values of $\xi = 2\eta$. Moreover, every value of ξ is associated with a multiplier θ that is greater than one. Finally, we can exchange the order of the min and max, which implies that $u = 0, f = f_{\theta}$ is a Nash equilibrium as well, where θ is chosen to satisfy the entropy constraint for a given value of ξ .

Thus, by expanding the set of admissible perturbations from mean shifts to arbitrary (absolutely continuous) density shifts, we have been able to avoid some of the complications of game (5.4.1). But we continue to be led to study limiting decision rules as θ decreases to some critical value, namely $\theta = 1$ in this example. The breakdown point for θ will no longer be associated with a finite value of ξ . The limiting solution as $\theta \searrow 1$ corresponds to the H_{∞} control in chapter 7.

Introducing a translation term b into the objective as in

$$-\frac{u^2}{2} - \frac{\int \left(u - b + \epsilon\right)^2 f\left(\epsilon\right) d\epsilon}{2}$$

will cause the worst-case distribution to have a nonzero mean, but there will still be a variance enhancement. The quadratic objective makes the worst-case distribution remain normal. The enhanced variance will not alter the decision for u. Thus the multiplier solution for u in (5.5.1) also solves the stochastic game (5.7.3). However, the implied variance enhancement is needed to match multipliers and constraints for the stochastic game.

5.8. Concluding remarks

This chapter has displayed two types of zero-sum two-player games that induces decisions that are robust to model misspecification. Each game has a malevolent nature choose a model misspecification to frustrate the decision maker. The 'constraint game' directly constrains the distortions to the approximating model that the malevolent agent can make. The 'multiplier game' penalizes those distortions. The two games are equivalent under conditions that allow us to invoke the Lagrange multiplier theorem. For our simple static example, we displayed conditions under which the two games are equivalent, and explored conditions under which they capture concern for model misspecification.

We have considered two classes of misspecifications, one that allows distortions only in the mean of a Gaussian random variable, the other than allows arbitrary alternative density functions that satisfy a constraint on entropy. In the static setting of this chapter, for the first class of mean misspecifications only, misspecification is confined to not knowing the mean of a random shock or a constant term in a linear equation. Subsequent chapters take up models where the decision maker fears misspecified dynamics, which he expresses by allowing a distortion w to be the conditional mean of a shock vector. By allowing that conditional mean to feed back on the history of the state, a variety of misspecifications can be modelled. Thus, the following two chapters return to our main theme of dynamic games that can be used to design robust decision rules. The conceptual issues connecting the constraint game and the multiplier game will carry over to the richer setting of chapters 6 and 7.

A. Rational expectations equilibrium

The Phillips curve example of section 5.2 took π_e as given. This appendix constructs a rational expectations version of the model and shows how to compute a time-consistent or Nash equilibrium rate of inflation. We proceed by adapting some concepts of Stokey (1989) to this example. Thus, we define a Nash equilibrium (with robustness) for the model as follows:

Definition 5.A.1. Given multiplier $\theta > 1$, a *Nash equilibrium* is a pair (π, π_e) such that (a) $\pi = B(\pi_e; \theta)$, and (b) $\pi = \pi_e$. Here *B* is the government's best response map (5.2.11).

Condition (a) says that given π_e , the government is choosing a robust rule associated with multiplier θ . Condition (b) imposes rational expectations. It is easy to compute a rational expectations equilibrium by solving (5.2.7) and $\pi = \pi_e$ for π_e :

$$\pi_e(\theta) = \frac{\theta}{\theta - 1} U^* \gamma.$$
 (5.A.1)

Notice that $\pi'_e(\theta) < 0$, $\lim_{\theta \nearrow \infty} \pi_e(\theta) = U^*\gamma$, and $\lim_{\theta \searrow 1} \pi_e(\theta) = +\infty$. If the approximating model is true, so that the government's concern about misspecification is misplaced, the government's ignorance of the model causes it to set inflation higher than if it knew the model for sure.

Notice that Definition 5.A.1 imputes a concern for model misspecification to the government, but not to the private forecasters, who are assumed to know the π chosen by the government. In chapter 16 we shall return to discuss an alternative version of rational expectations that imposes more symmetry between the government and private agents.

Chapter 6 Time domain games for robustness

6.1. Several games have identical outcomes

This chapter studies two-player zero-sum dynamic games in which a minimizing player helps a maximizing player design a decision rule that is robust to misspecification of a dynamic approximating model that links controls today to state variables tomorrow. We represent misspecification by allowing shocks to feed back on the history of the state in ways that an approximating model excludes. We generalize the static constraint and multiplier games of chapter 5 to this dynamic setting. The constraint and multiplier games differ in how they parameterize a set of alternative specifications that surround an approximating model. The constraint games require that the entropy of each alternative model relative to the approximating model not exceed a nonnegative parameter η . The multiplier games restrict relative entropy implicitly via a penalty parameter θ . If the parameters η and θ are appropriately related, the constraint and multiplier games have equivalent outcomes.

We begin the chapter by studying three multiplier games that have identical players, payoffs, and actions but different timing protocols. Games with different timing protocols usually have different outcomes, but because the two players' preferences are perfectly *mis*aligned, our games have identical outcomes. We devote much of this chapter to verifying the equivalence of outcomes and equilibrium representations of multiplier games for our three timing protocols. After that, we show how the equilibrium of a multiplier game provides an equilibrium of a constraint game.

The equivalence of outcomes of multiplier games across different timing protocols is of substantial importance. We shall exploit it frequently in the economic applications in subsequent chapters, for example, in the equilibrium in a model with a Ramsey planner that we propose in chapter 16. While some of the proofs in this chapter involve complicated arguments, they justify simple

- 138 -

algorithms and appealing ways of interpreting robust decision rules. We summarize these compactly in section 6.2.2. The reader interested only in results can extract practical procedures from this section.

6.2. The setting

A decision maker has a unique explicitly specified approximating model but concedes that the data might actually be generated by a member of a set of models that surround the approximating model. One parameter, either θ or η , measures a set of perturbations to the approximating model. Three models within the set are especially important: the decision maker's approximating model; an unknown true model that generates the data; and a worst case model that emerges as a by-product of a robust decision making procedure. Each model specifies that an $n \times 1$ state vector evolves according to

$$x_{t+1} = Ax_t + Bu_t + Cw_{t+1} \tag{6.2.1}$$

where x_0 is given, u_t is a vector of controls, and w_{t+1} is a vector of specification errors. The approximating model assumes that $w_t = 0 \ \forall t \ge 1$. The other models have $w_t \ne 0$ for some $t \ge 1$.¹ We assume that the matrix A has all of its eigenvalues inside the circle Γ in the complex plane, where $\Gamma = \{\zeta : |\zeta| = \frac{1}{\sqrt{\beta}}\}$.² Under this restriction on the eigenvalues of A,

$$(I - \zeta A)^{-1} = \sum_{j=0}^{\infty} A^j \zeta^j$$

¹ In chapter 8, we will consider stochastic models formed by replacing w_{t+1} by the sum of an i.i.d. Gaussian vector $\tilde{\epsilon}_{t+1}$ with mean zero and identity covariance matrix I and a distortion w_{t+1} that is measurable with respect to the history of x_t . The presence of ϵ_{t+1} obscures the model misspecification with noise. This setting lets us use model detection error probabilities to calibrate the value of θ .

² A more general but still workable assumption is that the pair $(\sqrt{\beta}A, B)$ is *stabilizable*, where $\beta \in (0, 1]$ is a discount factor. The pair $(\sqrt{\beta}A, B)$ is said to be stabilizable if there exists a matrix \tilde{F} for which $A - B\tilde{F}$ has all of its eigenvalues inside Γ . (See chapter 3, page 61 for more about stabilizability.) Under this condition, we can rewrite the system $x_{t+1} = Ax_t + Bu_t$ as $x_{t+1} = (A - B\tilde{F})x_t + B\tilde{u}_t$, where $u_t = -\tilde{F}x_t + \tilde{u}_t$, and then proceed to view \tilde{u}_t as the control.

is a convergent power series for ζ inside the circle Γ .

Define a target vector

$$z_t = Hx_t + Ju_t. ag{6.2.2}$$

The decision maker wants to maximize the objective function

$$-\sum_{t=0}^{\infty}\beta^t z_t' z_t. \tag{6.2.3}$$

We use the following measure of model misspecification:

$$R(w) = \sum_{t=0}^{\infty} \beta^{t+1} w'_{t+1} w_{t+1}.$$
(6.2.4)

Hansen, Sargent, Turmuhambetova, and Williams (2005XXX) refer to R(w) as entropy.

The decision maker believes that the data are generated by a model that satisfies $R(w) \leq \eta$ but is otherwise ignorant about $\{w_{t+1}\}$. The decision maker wants a decision rule that works well for any model satisfying $R(w) \leq \eta$.

6.2.1. Constraint and multiplier problems

Let u denote the sequence $\{u_t\}_{t=0}^{\infty}$ and w the sequence $\{w_{t+1}\}_{t=0}^{\infty}$. Two types of games induce robust decisions. First, for $\eta \in \Upsilon = \{\eta : 0 \le \eta \le \overline{\eta}\}$, we have:

Definition 6.2.1. The constraint robust control problem is

$$\sup_{u} \inf_{w} -\sum_{t=0}^{\infty} \beta^{t} z_{t}^{\prime} z_{t}$$
(6.2.5)

subject to (6.2.1) and $R(w) \leq \eta$.

Second, for θ belonging to a set $\Theta = \{\theta : 0 < \theta \le +\infty\}$, we define

Definition 6.2.2. The multiplier robust control problem is

$$\sup_{u} \inf_{w} - \left[\sum_{t=0}^{\infty} \beta^{t} z_{t}' z_{t} - \theta R\left(w\right)\right]$$
(6.2.6)

subject to (6.2.1).

The Lagrange multiplier theorem (see Luenberger (1969), pp. 216-221) connects the solutions of these two problems, as we shall discuss in detail in chapter 7. In equation (6.2.9) below, we shall describe how to determine the lower bound $\underline{\theta}$. The upper bound $\overline{\eta}$ is connected to $\underline{\theta}$ in a way that we describe in section 6.8. If θ and η are appropriately related, the multiplier and constraint problems are equivalent.

We begin by focusing on three versions of the multiplier problem (6.2.6). Section 6.8 discusses the relation between the constraint and multiplier problems and shows how the same pair of decision rules solves both problems.³ Section 6.9 formulates a recursive version of the constraint problem and links the derivative of its value function to the multiplier θ .

6.2.2. Operators and decision rules

It is convenient to summarize some operators that will occur frequently in this chapter and the next. There exists a lower bound or breakdown value $\underline{\theta}$ that describes the largest set of perturbed models against which it is possible to acquire robustness. After we have introduced several operators, we display condition (6.2.9), which provides a check for whether $\theta > \underline{\theta}$.⁴ For a given $\theta > \underline{\theta}$, we can compute robust linear decision rules $u_t = -F^*x_t, w_{t+1} = K^*x_t$ by using the following operators:

$$T(P) = H'H - H'J(J'J)^{-1}J'H + \beta \left[A' - H'J(J'J)^{-1}B'\right] \left[P - \beta PB(J'J + \beta B'PB)^{-1}B'P\right] \left[A - B(J'J)^{-1}J'H\right] (6.2.7a) or$$

$$T(P) = H'H + \beta A'PA - (\beta B'PA + H'J) \times (J'J + \beta B'PB)^{-1} (\beta B'PA + J'H)$$
(6.2.7b)

$$\mathcal{D}(P) = P + PC \left(\theta I - C'PC\right)^{-1} C'P \qquad (6.2.7c)$$

$$F(P) = (J'J + \beta B'PB)^{-1} (\beta B'PA + J'H).$$
(6.2.7d)

 $^{^3}$ Hansen, Sargent, Turmuhambetova, and Williams (2001) connect these problems in the context of an approximating model that is a continuous time diffusion.

⁴ Chapter 7 discusses $\underline{\theta}$ and the associated upper bound $\overline{\eta}$ that assure that the problems multiplier and constraint problems have finite values.

$$K(P) = (\theta I - C'PC)^{-1}C'P(A - BF(\mathcal{D}(P)))$$
(6.2.7e)

$$S(P) = H'_F H_F + \beta A'_F \mathcal{D}(P) A_F, \qquad (6.2.7f)$$

where $A_F = A - BF$ and $H_F = H - JF$. Define the fixed points of the following two algebraic Riccati equations:

$$\bar{P} = T\left(\bar{P}\right) \tag{6.2.8a}$$

$$P^* = T \circ \mathcal{D}\left(P^*\right). \tag{6.2.8b}$$

The decision rule without a concern about robustness is $u_t = F(\bar{P})x_t$, while for a given θ satisfying $\underline{\theta} \leq \theta \leq +\infty$, the decision rule with a concern for robustness is $u_t = -F(\mathcal{D}(P^*))x_t$. The worst case shock process associated with the robust decision rule is $w_{t+1} = K(P^*)x_t$.

If we take a fixed point $P^* = T \circ \mathcal{D}(P^*)$, we can verify that $\theta > \underline{\theta}$ by checking that

$$\log \det \left(\theta I - C' P^* C\right) > -\infty \tag{6.2.9}$$

or equivalently that the eigenvalues of $(\theta I - C'P^*C)$ are all positive.⁵ This follows from Theorem 7.5.4. Of course, this check requires that we can compute a fixed point of $T \circ \mathcal{D}$, which might not be possible for $\theta < \underline{\theta}$. An alternative and in a sense more practical way to assure that $\theta > \underline{\theta}$ is to check the condition

$$\log \det \left(\theta I - C' P_j C\right) > -\infty \tag{6.2.10}$$

for each iterate $P_j, j \ge 1$, where P_j is computed as $P_{j+1} = T \circ \mathcal{D}(P_j)$ starting from $P_0 = 0$.

The $T \circ \mathcal{D}$ operator can be calculated in one step as:

$$T \circ \mathcal{D}(P) = H'H - H'J(J'J)^{-1}J'H + \beta \begin{bmatrix} A' - H'J(J'J)^{-1}B' \end{bmatrix} \times \begin{bmatrix} P - \beta P(B \quad C) \begin{pmatrix} J'J + \beta P'BP & \beta B'PC \\ \beta C'PB & -\beta\theta I + \beta C'PC \end{pmatrix}^{-1} \begin{pmatrix} B' \\ C' \end{pmatrix} P \end{bmatrix} \begin{bmatrix} A - B(J'J)^{-1}J'H \end{bmatrix}.$$
(6.2.11)

For a given θ , a policy improvement algorithm for computing a robust decision rule iterates on the operators S and F:

 $^{^{5}}$ Chapter 7 calls the object on the left side of (6.2.9) 'entropy' and interprets it is a criterion that inspires a decision player to choose a decision rule that is robust to misspecification.

- 1. For a fixed decision rule F, define the associated operator S, and compute the fixed point P = S(P).
- 2. Compute a new decision rule F = F(P).
- 3. Iterate to convergence on steps 1 and 2.

Step 1 computes a value function attained by using a fixed decision rule F forever, where the 'distortion operator' \mathcal{D} evaluates future utilities. Step 2 finds an F that solves a two-period optimum problem, with $\mathcal{D}(P)$ being used to form the continuation value function. This is an efficient algorithm for computing a robust rule.⁶

By providing simple algorithms for computing a robust decision rule, formulas (6.2.7) and (6.2.8) completely summarize the practical content of this chapter. We will use these operators in subsequent sections to justify the equivalence of outcomes from distinct two-player zero-sum dynamic games. We shall say more about the *S* operator in chapter 7.

6.3. Timing protocols in three versions of a multiplier game

Timing protocols for three versions of the zero-sum two-player multiplier game (6.2.6) translate into differences in the spaces to which the maximizing player's choice of the sequence for u and the minimizing player's choice of the sequence for w are confined. We restrict u and w to one of the following spaces:

$$W = \{w : \sum_{t=1}^{\infty} \beta^{t} w'_{t} w_{t} < +\infty\}$$
$$U = \{u : \sum_{t=0}^{\infty} \beta^{t} u'_{t} u_{t} < +\infty\}$$
$$W_{K} = \{w : w_{t+1} = K x_{t}\}$$
$$U_{F} = \{u : u_{t} = -F x_{t}\}.$$

By choosing from among these sets, we create three versions of problem (6.2.6):

 $^{^{6}\,}$ Other efficient algorithms use a doubling algorithm to compute the fixed point of $T\circ\mathcal{D}\,.$ See chapter 3.

Game 1 (SEQ): a static multiplier game where both players choose sequences: $u \in U, w \in W$.

Game 2 (STACK): a Stackelberg multiplier game where the minimizing player chooses a sequence but the maximizing player chooses a rule that feeds back on the state: $u \in U_F, w \in W$.

Game 3 (MARKOV): a Markov perfect multiplier game where both players choose rules that feed back on the state: $u \in U_F, w \in W_K$.

We apply the same standard equilibrium concept to each game.

Definition 6.3.1. A Nash equilibrium is a pair w^*, u^* such that, given u^* , w^* is optimal for the *w*-choosing player and, given w^* , u^* is optimal for the *u*-choosing player.

6.4. Score card: timing and uses of three multiplier games

For future reference, it is useful to present Table 6.5.1. Within equilibria of each of our three multiplier games, Table 6.5.1 summarizes the different assumptions about what the *u*-choosing player and the *w*-choosing player regard as exogenous. Along an equilibrium path, choices can be represented as $u_t = -F^*x_t, w_{t+1} = K^*x_t$ with the same (F^*, K^*) in all three of our games. In game 1, neither player takes the feedback rule of the other player into account and instead regards the other player's sequence of choices as given. But in games 2 and 3, either one or both of the players takes the other player's feedback rule into account. A player that recognizes the other player's feedback rule can influence the parts of the other player's future decisions that feed back on future states.

The right column of Table 6.5.1 reveals much about the structure of players' optimum problems in the three games. That game 1 confronts each player with an arbitrary sequence of actions requires that we formulate each player's problem in the space of sequences, obtain Euler equations for both players, then compute an equilibrium by solving the system formed by stacking the two players' Euler equations. In game 1, neither player's problem is recursive in the state vector x_t . In game 2, the *w*-player's problem is recursive in the state vector x_t , while

in game 3, both players' problems are recursive in x_t . These recursive problems are evidently dynamic programming problems, while both players' problems in game 1 are not dynamic programming problems in the physical state vector x_t . However, there does exist a recursive representation of both players' problems in an equilibrium of game 1: it entails augmenting the physical state vector in a way that captures the limited power that both players have to influence the other player's sequence of choices in game 1.

Each of our three multiplier games has special uses: either for justifying a particular interpretation of a robust rule, or for bringing to light a particular aspect of the structure of a robust control problem, or for computing a robust decision rule.

- 1. Because it makes both players choose best responses to sequences, game 1 justifies characterizing the equilibrium decisions of the maximizing and minimizing players by stacking their Euler equations. We take advantage of this structure in chapter 16 when we pose and solve robust versions of Ramsey or Stackelberg problems. In these problems, a robust Stackelberg leader once and for all commits himself to a sequence of decisions taking into account the robust response of a group of followers, each of whom chooses sequentially. Off the equilibrium path of game 1, the decision problem of neither the maximizing nor the minimizing players in game 1 is recursive. However, by expanding the state space to include variables that allow a recursive representation of the equilibrium sequence chosen by the minimizing player, we can give a recursive representation for the problem of the maximizing player in game 1. This representation has two important uses. First, it justifies a 'Bayesian' interpretation of a robust control problem. That is, it displays a law of motion for forcing variables that is distorted relative to the approximating model and for which the robust decision rule is actually an optimal (non robust) decision rule. Second, this recursive representation of the maximizing player's problem reveals a certainty equivalence principle that applies to the maximizing player's decision.
- 2. By making the minimizing player a Stackelberg leader who chooses a sequence once and for all, game 2 becomes the natural setting for studying frequency domain representations of a robust control problem. We develop

Game	Player	Exogenous	Constraint
1	u	$\{w_{t+1}^*\}$	$x_{t+1} = Ax_t + Bu_t + Cw_{t+1}^*$
	w	$\{u_t^*\}$	$x_{t+1} = Ax_t + Bu_t^* + Cw_{t+1}$
2	u	$\{w_{t+1}^*\}$	$x_{t+1} = Ax_t + Bu_t + Cw_{t+1}^*$
	w	$u_t = -F^* x_t$	$x_{t+1} = (A - BF^*)x_t + Cw_{t+1}$
3	u	$w_{t+1} = K^* x_t$	$x_{t+1} = (A + K^*C)x_t + Bu_t$
	w	$u_t = -F^* x_t$	$x_{t+1} = (A - BF^*)x_t + Cw_{t+1}$

Table 6.5.1: Players' constraints in three multiplier games

 H_{∞} and so-called minimum entropy representations of robust control problems in chapter 7, both of which are cast in the frequency domain and rest directly on game 2.

3. The Markov perfect game 3 justifies a very useful algorithm that computes a robust decision rule and the matrix P in the quadratic form for the associated value function by iterating to convergence of $T \circ \mathcal{D}$. Here the T operator summarizes the maximizing player's choice and the \mathcal{D} operator summarizes the minimizing player's choice in their 'two-period' problems.

6.5. Game 1: The multiplier game in sequences

In game 1 (SEQ), the objective of the two players can be written:

$$\mathcal{C} = \sum_{t=0}^{\infty} \beta^t \left(-z'_t z_t + \beta \theta w'_{t+1} w_{t+1} \right), \qquad (6.5.1)$$

subject to the state-evolution equation (6.2.1) and the target vector relation (6.2.2). The initial state vector x_0 is given. A maximizing player chooses $u \in U$ and a minimizing player chooses $w \in W$.

Definition 6.5.1. An equilibrium of the multiplier game in sequences (SEQ) is a pair of sequences $u^* \in U, w^* \in W$ that solve both players' problems.

To make the minimization problem well posed, we must restrict the value of θ . The penalty term $\theta w'_{t+1}w_{t+1}$ is convex in w_{t+1} by construction, but the objective function contains an additional contribution from w_{t+1} because w_{t+1} alters subsequent targets and the objective is *concave* in these targets. For a sufficiently large penalty parameter θ , the first term dominates and this assures the convexity of the entire intertemporal objective in w. A condition that assures convexity is that θ is sufficiently large that $\theta I - G'G$ is positive definite on $\Gamma = \{\zeta : |\zeta| = \frac{1}{\sqrt{\beta}}\}$, where $G = H(I - \zeta A)^{-1}C$. We will motivate this restriction on θ further in chapter 7.

To find an equilibrium of game SEQ, we begin by substituting from (6.2.2) for z_t in the objective function and then form a Lagrangian for each player. The zero-sum objective function implies that these two Lagrangians have first-order conditions that impart identical laws of motion to the two players' co-state variables; the initial values of the costate variables must also be equal because of the common value function. These features allow us to analyze the game by forming a *single* Lagrangian:

$$\mathcal{L} = -\sum_{t=0}^{\infty} \beta^{t} \{ (x'_{t}H'Hx_{t} + u'_{t}J'Ju_{t} + 2u'_{t}J'Hx_{t}) \\ 2\beta\mu'_{t+1} (Ax_{t} + Bu_{t} + Cw_{t+1} - x_{t+1}) - \beta\theta w'_{t+1}w_{t+1} \}.$$
(6.5.2)

We proceed by using the first-order conditions from this single Lagrangian to obtain a candidate equilibrium, then verify that this candidate equilibrium is indeed the outcome of the two constrained optimization problems that appear in Definition 6.5.1.

To generate a candidate equilibrium, notice that first-order conditions for the Lagrangian with respect to u_t, w_{t+1}, x_{t+1} , respectively, are:

$$J'Ju_t + J'Hx_t + \beta B'\mu_{t+1} = 0$$

- $\theta w_{t+1} + C'\mu_{t+1} = 0$
 $\beta A'\mu_{t+1} + H'Hx_t + H'Ju_t - \mu_t = 0.$ (6.5.3)

Assume that J'J is nonsingular and solve for u_t and w_{t+1} :

$$u_{t} = -(J'J)^{-1} J'Hx_{t} - \beta (J'J)^{-1} B'\mu_{t+1}$$
(6.5.4)

$$w_{t+1} = \frac{1}{\theta} C' \mu_{t+1}. \tag{6.5.5}$$

Substitute these expressions for u_t and w_{t+1} into the state equation to get

$$x_{t+1} = \left[A - B(J'J)^{-1}J'H\right]x_t - \left[\beta B(J'J)^{-1}B' - \frac{1}{\theta}CC'\right]\mu_{t+1}.$$

Substituting the same expressions into (6.5.3) gives

$$\beta \left[A' - H'J (J'J)^{-1} B' \right] \mu_{t+1} + \left[H'H - H'J (J'J)^{-1} J'H \right] x_t - \mu_t = 0.$$

Write the system as

 $L\begin{bmatrix} x_{t+1}\\ \mu_{t+1} \end{bmatrix} = N\begin{bmatrix} x_t\\ \mu_t \end{bmatrix}$ (6.5.6)

where

$$L = \begin{pmatrix} I & \left[\beta B \left(J'J\right)^{-1} B' - \frac{1}{\theta} C C'\right] \\ 0 & \beta \left[A' - H'J \left(J'J\right)^{-1} B'\right] \end{pmatrix}$$

and

$$N = \begin{pmatrix} \left[A - B (J'J)^{-1} J'H \right] & 0 \\ - \left[H'H - H'J (J'J)^{-1} J'H \right] & I \end{pmatrix}.$$

It can be verified that the matrix pencil $(\frac{\lambda}{\sqrt{\beta}}L - N)$ is symplectic.⁷ It follows that the generalized eigenvalues of (L, N) come in $\sqrt{\beta}$ -symmetric pairs: for every eigenvalue λ_i , there is another eigenvalue λ_{-i} such that $\lambda_i \lambda_{-i} = \beta^{-1}$.

To assure existence of a candidate equilibrium, we rule out generalized eigenvalues of (L, N) on the circle $\Gamma = \{\zeta : |\zeta| = \frac{1}{\sqrt{\beta}}\}$ so that half of the generalized eigenvalues are inside the circle Γ and the other half are outside this circle. The generalized eigenvectors associated with the eigenvalues inside Γ generate the $(\sqrt{\beta})$ -stable deflating subspace. The dimension of this subspace equals the number of entries in the state vector x_t . We assume that there exists a positive semidefinite matrix P^* such that the stable deflating subspace can be represented as $\begin{pmatrix} I \\ P^* \end{pmatrix} x$. Under these restrictions, we can construct a candidate equilibrium with $\mu_t = P^* x_t$ and a state vector sequence that satisfies

$$L\begin{pmatrix}I\\P^*\end{pmatrix}x_{t+1} = N\begin{pmatrix}I\\P^*\end{pmatrix}x_t.$$
(6.5.7)

That the candidate equilibrium is indeed an equilibrium can be verified under conditions that we summarize in:

Theorem 6.5.1. Suppose that

(i) (A, B) is stabilizable.

⁷ See chapter 3 for the definition and properties of symplectic pencils.

- (ii) (H, A) is detectable⁸ and J'J is nonsingular.
- (iii) $\theta I G'G$ is positive definite on Γ where $G = H(I \zeta A)^{-1}C$.
- (iv) (L, N) has no generalized eigenvalues on Γ .

(v) an element of the $(\sqrt{\beta})$ stable deflating subspace of (L, N) can be represented as $\begin{pmatrix} I \\ P^* \end{pmatrix} x$ for some vector x and a given matrix P^* . Then there exist K^* and F^* for which an equilibrium of game SEQ is $u_t = -F^*A^{*t}x_0$ and $w_{t+1} = K^*A^{*t}x_0$, where $A^* = A - BF^* + CK^*$ has eigenvalues that are inside Γ . The matrix P^* is necessarily symmetric and the date zero value of the game is $-x_0'P^*x_0$. Also, $F^* = (J'J)^{-1}(J'H + \beta B'P^*A^*), K^* = \frac{1}{\theta}C'P^*A^*$.

Proof. We have already computed a candidate equilibrium by stacking the statecostate equations of the two players to get the linear difference equation system (6.5.7). The candidate equilibrium is a $\sqrt{\beta}$ stable sequence of state vectors that satisfies (6.5.7). Given conditions (iv) and (v), from the first partition of (6.5.7), we see that

$$\left(I + \left[\beta B \left(J'J\right)^{-1} B' - \frac{1}{\theta} C C'\right] P^*\right) x_{t+1} = \left[A - B \left(J'J\right)^{-1} J' H\right] x_t. \quad (6.5.8)$$

It follows from Theorem 21.7 of Zhou, Doyle and Glover (1996) that P^* is symmetric and that the matrix on the left side of (6.5.8) is nonsingular. Hence we have the state evolution:

$$x_{t+1} = A^* x_t$$

where

$$A^{*} = \left(I + \left[\beta B (J'J)^{-1} B' - \frac{1}{\theta} CC'\right] P^{*}\right)^{-1} \left[A - B (J'J)^{-1} J'H\right].$$

Using the same reasoning that led to equation (3.3.9), it can be shown that

$$\begin{pmatrix} I + \left[\beta B \left(J'J \right)^{-1} B' - \frac{1}{\theta} C C' \right] P^* \end{pmatrix}^{-1} \\ = I - \beta \left(B \quad C \right) \begin{pmatrix} J'J + \beta B'P^*B & \beta B'P^*C \\ \beta C'P^*B & -\beta \theta I + \beta C'P^*C \end{pmatrix}^{-1} \begin{pmatrix} B'P^* \\ C'P^* \end{pmatrix}.$$

⁸ Or equivalently (A', H') is stabilizable.

Therefore,

$$A^* = A - BF^* + CK^*$$

where F^* and K^* satisfy

$$F^* = (J'J)^{-1} (J'H + \beta B'P^*A^*)$$

$$K^* = \frac{1}{\theta} C'P^*A^*.$$
(6.5.9)

By (v), A^* has eigenvalues that are inside the circle Γ . Moreover, the first-order conditions from the Lagrangian (6.5.2) imply

$$w_{t+1} = K^* x_t$$

$$u_t = -F^* x_t.$$
(6.5.10)

More will have to be added here.

AAAAA Tom: we now reference all of the Assumptions, but I am not sure how we use the symmetry of P. I also need to check the Glover, Doyle, Zhou reference.

Conditions (i) and (ii) occur in the standard control theory summarized in chapter 3 and assure the existence of an optimal control that stabilizes the state in the absence of concerns about misspecification. In particular, they guarantee that the objective of the maximizing decision maker is strictly concave in the usequence. Condition (iii) guarantees that the objective is strictly convex in the w sequence.

6.5.1. Recursive formulation of maximizing player's problem

In game 1, each player chooses a sequence, taking as given the sequence chosen by the other player. Because an arbitrary sequence chosen by the other player does not have a recursive representation, out of equilibrium the problem of a player who must choose a best response to such an arbitrary *sequence* will not have a recursive representation. Nevertheless, Theorem 6.5.1 indicates that *in equilibrium* each player's sequence does have a recursive representation, the time t decision of each being a linear function of the state x_t as in (6.5.10). But (6.5.10) fails to embody the assumption that player *i*'s decisions do not

150

influence the sequence chosen by player -i, where -i means not i. For example, to represent the problem of the maximizing player in game 1, it won't work just to confront that player with $w_{t+1} = K^* x_t$ coupled with the law of motion $x_{t+1} = Ax_t + Bu_t + Cw_{t+1}$, because that would contradict the assumption that w is taken as an exogenous sequence by the maximizing (*u*-choosing) player.

The way around this problem is to introduce an additional state vector \hat{x}_t that satisfies the requirements that (1) \hat{x} cannot be controlled by the maximizing player and (2) it permits a recursive representation of the equilibrium choice of w. This can be accomplished by positing that

$$\hat{x}_{t+1} = A^* \hat{x}_t \tag{6.5.11a}$$

$$w_{t+1} = K^* \hat{x}_t \tag{6.5.11b}$$

where, as in Theorem 6.5.1, the matrix $A^* = A - BF^* + CK^*$. If we impose $\hat{x}_0 = x_0$, then (6.5.11) recursively generates the equilibrium $w \in W$ given by (6.5.10). Expressing w in terms of the uncontrollable state \hat{x}_t preserves the game 1 assumption that $w \in W$ is taken as given by the maximizing player; it also gives w a recursive representation and thereby allows us to apply dynamic programming to the maximizing player's problem.

We thus use representation (6.5.11) to pose a recursive version of the maximizing player's problem within the game 1 equilibrium. In particular, he maximizes (6.5.1) by choice of $\{u_t\}$ subject to the state evolution:

$$x_{t+1} = Ax_t + Bu_t + CK^* \hat{x}_t$$

$$\hat{x}_{t+1} = A^* \hat{x}_t$$

$$w_{t+1} = K^* \hat{x}_t.$$

(6.5.12)

Notice that u_t influences subsequent positions of x_s but not of \hat{x}_s and therefore not subsequent values of w_{s+1} . Equation (6.5.12) thereby captures the idea that the maximizing player takes the w sequence as given.

The optimizer of problem (6.5.12) is a decision rule $u_t = -\begin{bmatrix} \bar{F} & \hat{F} \end{bmatrix} \begin{bmatrix} x_t \\ \hat{x}_t \end{bmatrix}$ where $F^* = \bar{F} + \hat{F}$. Theorem 6.5.2 below verifies this and also that if we set $\hat{x}_0 = x_0$, then the outcome for this problem satisfies $\hat{x}_t = x_t$ for all $t \ge 1$.

The problem of maximizing (6.5.1) subject to (6.5.12) is obviously not useful for computing the F^* component of an equilibrium of game 1: to pose the problem, we must already know the equilibrium (K^*, F^*) that emerge in an equilibrium of game 1. However, the problem *is* useful as a tool for interpreting and decomposing the F^* component of game 1. Soon we shall discuss a 'Bayesian interpretation' of a robust decision rule that is based on the recursive version of the maximizing player's problem in an equilibrium of game 1. We shall use that interpretation extensively in our analysis of asset pricing in chapters 11 and 12. Before we discuss the Bayesian interpretation of this problem, we state a theorem about the value function and decision rule that solve that problem.

This theorem exploits the insight that the problem of maximizing (6.5.1) subject to (6.5.12) takes the form of what Anderson *et. al.* (1996) and chapter 3 call an *augmented regulator problem*. This allows us to break it into subproblems, the first of which is simply the non-robust ($\theta = +\infty$) version of the *u*-player's decision problem with $u_t = -\bar{F}x_t$ being the decision rule and $-x_0\bar{P}x_0$ being the value function for the ordinary control problem. These objects are constituents of the following:

Theorem 6.5.2. Consider an ordinary (non-robust) optimal linear regulator with current period objective

$$(Hx_t + Ju_t)'(Hx_t + Ju_t) - \beta\theta(K\hat{x}_t) \cdot (K\hat{x}_t)$$
(6.5.13)

subject to the law of motion

$$\begin{pmatrix} x_{t+1} \\ \hat{x}_{t+1} \end{pmatrix} = \begin{pmatrix} A & \hat{A} \\ 0 & A^* \end{pmatrix} \begin{pmatrix} x_t \\ \hat{x}_t \end{pmatrix} + \begin{pmatrix} B \\ 0 \end{pmatrix} u_t$$
(6.5.14)

where $\hat{A} = CK^*$ and $A^* = A - BF^* + CK^*$. Then the optimal value function is

$$-\begin{pmatrix} x_0\\ \hat{x}_0 \end{pmatrix}' \begin{pmatrix} \bar{P} & \hat{P}\\ \hat{P}' & \tilde{P} \end{pmatrix} \begin{pmatrix} x_0\\ \hat{x}_0 \end{pmatrix}.$$

where

$$\hat{P} = P^* - \bar{P}$$
$$\tilde{P} = \bar{P} - P^*$$

and where \overline{P} is the stabilizing solution to the Riccati equation for the ordinary (non-robust) control problem and P^* is the stabilizing solution to the Riccati equation for the robust control problem. The optimal control law is

$$u_t = -\bar{F}x_t - \hat{F}\hat{x}_t$$

where

$$\bar{F} = \left(J'J + \beta B'\bar{P}B\right)^{-1} \left(\beta B'\bar{P}A + J'H\right)$$
$$\hat{F} = \left(J'J + \beta B'\bar{P}B\right)^{-1} \left(\beta B'\bar{P}\hat{A} + \beta B'P^*A^*\right).$$

Moreover, $u_t = -\bar{F}x_t$ is the control law for the ordinary (non-robust) problem and $\hat{F} + \bar{F} = F^*$ where $u_t = -F^*x_t$ is the control law for the robust control problem.

Proof. The matrices \bar{P} and P^* are fixed points for the Riccati equations for the ordinary and robust linear regulators, respectively, so that $\bar{P} = T(\bar{P})$ and $P^* = T \circ \mathcal{D}(P^*)$. The proof proceeds by solving the augmented linear regulator defined by the problem (6.5.13), (6.5.14), which leads us to compute $\bar{P}, \hat{P}, \tilde{P}$ recursively; and then by verifying that these matrices solve the following equations $\bar{P} = T(\bar{P}), \ \bar{P} + \hat{P} = T \circ \mathcal{D}(\bar{P} + \hat{P}), \ \tilde{P} - \bar{P} = T \circ \mathcal{D}(\tilde{P} - \bar{P}),$

Because the optimization problem (6.5.13), (6.5.14) is an augmented linear regulator problem (see chapter 3), we can solve it in three steps. In the first step, we set $\hat{x}_0 = 0$. This makes the sequence \hat{x}_t disappear from the problem. Let \bar{P} denote the matrix that stabilizes the corresponding deflating subspace so that \bar{P} solves the algebraic Riccati equation $\bar{P} = T(\bar{P})$ or

$$\beta \left[A' - H'J (J'J)^{-1} B' \right] \left[P - \beta PB (J'J + \beta B'PB)^{-1} B'P \right] \left[A - B (J'J)^{-1} J'H \right] + H'H - H'J (J'J)^{-1} J'H = P.$$

Let \overline{F} denote the control law for the ordinary (non-robust) control problem given by:

$$\bar{F} = \left(J'J + \beta B'\bar{P}B\right)^{-1} \left(\beta B'\bar{P}A - J'H\right).$$

Define $\bar{A} = A - B\bar{F}$. The matrix \bar{P} also solves the Sylvester equation:

$$P = \left(H - J\bar{F}\right)' \left(H - J\bar{F}\right) + \beta \bar{A}' P\bar{A}.$$

In the second step, we activate the uncontrollable state \hat{x}_t and compute \hat{P} . The optimal control law is

$$u_t = -\bar{F}x_t - \hat{F}\hat{x}_t$$

and $P = \hat{P}$ solves the Sylvester equation:

$$\beta \left(A - B\bar{F} \right)' \left(\bar{P}\hat{A} + PA^* \right) = P.$$

Equivalently, $P = \hat{P}$ solves

$$\beta \left[A' - H'J (J'J)^{-1} B' \right] \left[\bar{P} - \beta \bar{P}B (J'J + \beta B' \bar{P}B)^{-1} B' \bar{P} \right] \hat{A} + \beta \left[A' - H'J (J'J)^{-1} B' \right] \left[P - \beta \bar{P}B (J'J + \beta B' \bar{P}B)^{-1} B' P \right] A^* = P.$$

The matrix \hat{P} that solves this Sylvester equation equals $\hat{P} = P^* - \bar{P}$ where P^* solves the Riccati equation $P^* = T \circ \mathcal{D}(P^*)$ that is associated with the robust control problem, which from (6.2.11) or (6.*B*.15) is

$$\beta \left[A' - H'J (J'J)^{-1} B' \right] \left[P^* - \beta P^* (B \ C) \begin{pmatrix} J'J + \beta B'P^*B & \beta B'P^*C \\ \beta C'P^*B & -\beta \theta I + \beta C'P^*C \end{pmatrix}^{-1} \begin{pmatrix} B' \\ C' \end{pmatrix} P^* \right] \\ \left[A - BJ (J'J)^{-1} J'H \right] + H'H - H'J (J'J)^{-1} J'H = P^*.$$

In appendix C, we verify that $\hat{P} = P^* - \bar{P}$. The portion of the control law that feeds back onto \hat{x} is

$$\hat{F} = \left(J'J + \beta B'\bar{P}B\right)^{-1} \left(\beta B'\bar{P}\hat{A} + \beta B'\hat{P}A^*\right).$$

In the third step, we compute \tilde{P} , which solves the Sylvester equation:

$$P = -\theta K^{*'} K^{*} + \beta A^{*'} P A^{*} + \beta \hat{A}' \left[\bar{P} - \beta \bar{P} B \left(J' J + \beta B' \bar{P} B \right)^{-1} B' \bar{P} \right] \hat{A} + \beta A^{*'} \left[\hat{P} - \beta \hat{P} B \left(J' J + \beta B' \bar{P} B \right)^{-1} B' \bar{P} \right] \hat{A} + \beta \hat{A}' \left[\hat{P} - \bar{P} B \left(J' J + \beta B' \bar{P} B \right)^{-1} B' \hat{P} \right] A^{*}.$$

6.5.2. 'Bayesian' interpretation of robust control

In Theorem 6.5.2, the maximizing player's decision problem takes the form of an ordinary (non-robust) control problem in which the law of motion for the state x is distorted by adding $CK^*\hat{x}_t$ to the right side of the law of motion for x_{t+1} in the approximating model. In stochastic versions of this decision problem, the existence of such a representation for the robust decision maker's decision problem assures us that there is a 'Bayesian' interpretation of a robust decision rule in the sense that there exists *some* law of motion for the augmented state (x_t, \hat{x}_t) for which the robust decision rule would be optimal. This establishes that the robust decision rule is 'undominated' and so can be said to be 'admissible' in the Bayesian sense.⁸

6.5.3. Recursive version of minimizing player's problem

We could proceed in a symmetric way to pose a recursive version of the minimizing player's problem in game 1. In particular, that player would minimize (6.5.1) by choice of $\{w_{t+1}\}$ subject to the state evolution:

$$x_{t+1} = Ax_t + Cw_{t+1} - BF^* \hat{x}_t$$

$$\hat{x}_{t+1} = A^* \hat{x}_t$$

$$u_t = -F^* \hat{x}_t.$$
(6.5.15)

This problem also takes the form of an *augmented regulator problem*. The minimizer is a decision rule $w_{t+1} = \begin{bmatrix} \bar{K} & \hat{K} \end{bmatrix} \begin{bmatrix} x_t \\ \hat{x}_t \end{bmatrix}$.

 $^{^{8}}$ Appendix B of chapter 13 gives a related Bayesian interpretation of a robust Kalman gain for a filtering problem.

6.6. Game 2: the Stackelberg multiplier game

We now turn to game 2 (STACK), which relative to game 1 has a timing protocol that gives the minimizing player more power to influence the choice of the maximizing player. We continue to withhold from the maximizing player any power to influence subsequent values of the minimizing player's decision.⁹

Section 6.5.3 posed a recursive version of the problem of the minimizing player within an equilibrium of game 1 (SEQ). As emphasized in section 6.5.1, what mattered for attaining a recursive representation of player i's problem in equilibrium is having a recursive representation of player not-i's decision sequence. In an equilibrium of game 2, the constraints facing the minimizing player have a recursive representation but differ substantially from (6.5.15). In game 2 (STACK), the maximizing player chooses sequentially and in equilibrium chooses a decision rule $u_t = -F^*x_t$ that feeds back on the state. Given the maximizing player's choice of feedback rule, the minimizing player chooses a sequence $w \in W$. Therefore, in equilibrium, the minimizing player faces the law of motion

$$x_{t+1} = Ax_t + Cw_{t+1} - BF^*x_t, (6.6.1)$$

which should be compared with the law of motion (6.5.15) that the minimizing player faces in the recursive version of game 1. Under (6.5.15), the w player regards the u sequence as fixed. But under (6.6.1), the w-setting player recognizes that he can influence future u's because $u_t = -F^*x_t$. Despite their different timing protocols, identical objects (K^*, F^*) characterize the equilibria of both games 2 (STACK) and 1 (SEQ).

Theorem 6.6.1. Let $A_F = A - BF$, $H_F = H - JF$, and

$$K(F) = (\theta I - C'PC)^{-1}C'P(A - BF), \qquad (6.6.2)$$

where P is the positive semidefinite solution to the Riccati equation

$$P = H'_{F}H_{F} + \beta A'_{F}PA_{F} + \beta A'_{F}PC (\theta I - C'PC)^{-1} C'PA_{F}$$
(6.6.3)

for which $A_F + CK$ has eigenvalues that are inside the circle Γ . Suppose (i) J'J is nonsingular and $J + H\zeta(I - \zeta A)^{-1}B$ has full column rank on Γ ;

 $^{^{9}}$ The maximizing player has to wait until game 3 to acquire such power.

(ii) The matrix pencil defined by (L, N) in (6.5.6) has no generalized eigenvalues on Γ ;

(iii) Any element of the $(\sqrt{\beta})$ deflating subspace of (L, N) can be represented as $\begin{pmatrix} I \\ P \end{pmatrix} x$ for some vector x where P is a symmetric, positive semidefinite matrix;

(iv) $\theta I - C'PC$ is positive definite.

Then there exists an equilibrium of the Stackelberg multiplier game in which $F = F^*$ and $K = K(F^*)$; consequently $w_{t+1}(F^*) = K^*(A^*)^t x_0$, where F^* , K^* and A^* are the same matrices that represent the equilibrium of game SEQ.

Conditions (ii) and (iii) are assured when $\beta B(J'J)^{-1}B' - \frac{1}{\theta}CC'$ is positive semidefinite. However, this positive semidefiniteness condition is much stronger than what is actually needed for many applications.

Proof. In the Stackelberg multiplier game, the maximizing player submits a decision rule $u_t = -Fx_t$. The minimizing player chooses a sequence $\{w_{t+1}(F)\}$ to minimize (6.5.1). For some F's, the infimum may not be attained. We can form the criterion $C(F, x_0)$, noting that it may be $-\infty$ for some choices of F. We wish to show that

$$C\left(F^*, x_0\right) \ge C\left(F, x_0\right)$$

for any $F \in \mathcal{F}$.

To verify this inequality, we first show that $\{w_{t+1}(F^*)\}$ coincides with $\{w_{t+1}^*\}$ of the equilibrium of game SEQ. Thus, we study the problem of minimizing (6.5.1) by choice of $w \in W$ subject to

$$x_{t+1} = (A - BF^*) x_t + Cw_{t+1}.$$
(6.6.4)

This differs from the optimum problem of the malevolent agent (over w) within a game 1 SEQ equilibrium because now the malevolent agent does not regard the control sequence u as an exogenous element of W but instead knows that u_t feeds back on the state via $u_t = -F^*x_t$, a description that is embedded in (6.6.4), and therefore does not carry along \hat{x}_t as a separate component of the state vector as he did in (6.5.12) for game STACK. Here the malevolent agent knows that $x_t = \hat{x}_t$ when solving his optimization problem.

To show that the minimizing $\{w_{t+1}(F^*)\}$ coincides with $\{w_{t+1}^*\}$ from the SEQ equilibrium, we form the discrete-time Hamiltonian system for choosing

 $\{w_{t+1}\}\$ as a function of F^* , as required in the Stackelberg multiplier equilibrium. Recalling that $H_F = H - JF$, from (6.5.3) the first-order conditions for w_{t+1} collapse to:

$$-\theta w_{t+1} + C' \mu_{t+1} = 0$$

$$\beta A_o' \mu_{t+1} + H_F' H_F x_t - \mu_t = 0.$$

check the above for H_F versus H. Next impose that $u_t = -F^*x_t$ and let $H_F^* = H - JF^*$ and $A_F^* = A - BF^*$. Then from (6.5.4) $BF^*x_t = -B(J'J)^{-1}J'H^*x_t - \beta B(J'J)^{-1}B'\mu_{t+1}$. The state equation becomes $x_{t+1} = A_F^*x_t + \theta^{-1}(C'C)\mu_{t+1}$. Note also from (6.5.4) that

$$-F^{*'}J'Ju_t = F^{*'}J'H_Fx_t + \beta F^{*'}B'\mu_{t+1}.$$

The modified co-state equation becomes $\beta A_F^{*\prime} \mu_{t+1} + H_F^{*\prime} H_F^* x_t - \mu_t = 0$, so that

$$\begin{pmatrix} I & -\frac{1}{\theta}CC'\\ 0 & \beta A_F^{*\prime} \end{pmatrix} \begin{pmatrix} I\\ P \end{pmatrix} x_{t+1} = \begin{pmatrix} A_F^* & 0\\ -H_F^*\prime H_F^* & I \end{pmatrix} \begin{pmatrix} I\\ P \end{pmatrix} x_t.$$
(6.6.5)

It follows that P satisfies Riccati equation

$$\beta A_F^{*\prime} P \left(I - \frac{1}{\theta} C C' P \right)^{-1} A_F^* - P + H_F^{*\prime} H_F^* = 0$$
 (6.6.6)

and therefore also satisfies¹⁰

$$P = H_F^{*'}H_F^* + \beta A_F^{*'}PA_F^* + \beta A_F^{*'}PC (\theta I - C'PC)^{-1} C'PA_F^*.$$

This value of P gives the unique solution in (6.6.5) that implies that the state vector sequence is $\sqrt{\beta}$ stable. From the proof of Theorem 6.5.1 it follows that $K = K^*$ and the positive definiteness of $\theta I - G'G$ follows from the restriction that $\theta I - C'PC$ is positive definite. From this result, we can compute $C(F^*, x_0)$ by simply evaluating the objective in game SEQ.

Now evaluate $C(F, x_0)$ for some other choice of F in \mathcal{F} . We can bound this criterion as follows. First, recursively generate the game 1 SEQ equilibrium $\{w_{t+1}^*\}$ sequence as

$$\hat{x}_{t+1} = A^* \hat{x}_t$$
$$w_{t+1}^* = K^* \hat{x}_t$$

¹⁰ This is verified in the proof of Theorem 7.5.4.

where $\hat{x}_0 = x_0$. Then form the state equation

$$x_{t+1} = (A - BF) x_t + CK^* \hat{x}_t$$
$$z_t = (H - JF) x_t.$$

Do we want K or K^* in the first equation? Using these recursions to evaluate (6.5.1), we obtain an upper bound $\hat{C}(F, x_0)$ on $C(F, x_0)$.

A convenient feature of this upper bound is that we can dominate $\hat{C}(F, x_0)$ by solving the following augmented regulator problem: maximize (6.5.1) by choice of a stabilizing control sequence $\{u_t\}$ for the state evolution

$$x_{t+1} = Ax_t + Bu_t + CK^* \hat{x}_t$$
$$\hat{x}_{t+1} = A^* \hat{x}_t$$

with $w_{t+1}^* = K^* \hat{x}_t$. But this is just the problem of the player who sets u_t in game SEQ. As in chapter 3, we solve this problem by stacking a state-costate system with the composite state (x_t, \hat{x}_t) and the costate corresponding to x_t . The costate for \hat{x}_t can be omitted because \hat{x}_t is an uncontrollable state vector. Thus we form a system

$$L^{a}\begin{pmatrix} x_{t+1}\\ \mu_{t+1}\\ \hat{x}_{t+1} \end{pmatrix} = N^{a}\begin{pmatrix} x_{t}\\ \mu_{t}\\ \hat{x}_{t} \end{pmatrix}$$

where:

$$L^{a} = \begin{pmatrix} I & \beta B (J'J)^{-1} B' & 0 \\ 0 & \beta \left[A' - H'J (J'J)^{-1} B' \right] & 0 \\ 0 & 0 & I \end{pmatrix}$$
$$N^{a} = \begin{pmatrix} A - B (J'J)^{-1} J'H & 0 & CK^{*} \\ - \left[H'H - H'J (J'J)^{-1} JH \right] & I & 0 \\ 0 & 0 & A^{*} \end{pmatrix}$$

•

To solve the problem we now look for the $\sqrt{\beta}$ deflating subspace of (L^a, N^a) parameterized as

$$\begin{pmatrix} x\\ P_2x + \hat{P}\hat{x}\\ \hat{x} \end{pmatrix} = \begin{pmatrix} I\\ P_2\\ 0 \end{pmatrix} (x - \hat{x}) + \begin{pmatrix} I\\ (P_2 + \hat{P})\\ I \end{pmatrix} \hat{x}.$$

We can simplify the problem to that of locating two $\sqrt{\beta}$ deflating subspaces of reduced dimension. The first deflating subspace is for the pair (L_2, N_2) with

$$L_{2} = \begin{pmatrix} I & \beta B (J'J)^{-1} B' \\ 0 & \beta \left[A' - H'J (J'J)^{-1} B' \right] \end{pmatrix}$$
$$N_{2} = \begin{pmatrix} A - B (J'J)^{-1} J'H & 0 \\ - \left[H'H - H'J (J'J)^{-1} JH \right] & I \end{pmatrix}.$$

This is the subspace associated with the component of $x_t - \hat{x}_t$ that must be set to zero to solve the control problem. Notice that (L_2, N_2) defines the statecostate system for the ordinary (non-robust) control problem. Thus we can restrict $x_t - \hat{x}_t$ to reside in the $\sqrt{\beta}$ stable deflating subspace of (L_2, N_2) using the matrix P for the ordinary control problem.

To study the second subspace, we seek a solution to:

$$L^{a}\begin{pmatrix} I\\ \left(P_{2}+\hat{P}\right)\\ I \end{pmatrix}\hat{x}_{t+1} = N^{a}\begin{pmatrix} I\\ \left(P_{2}+\hat{P}\right)\\ I \end{pmatrix}\hat{x}_{t}.$$

It is more convenient to pose this problem as being (a) to find a matrix \hat{P} such that we can represent the $\sqrt{\beta}$ deflating subspace of (\hat{L}, \hat{N}) as parameterized by:

$$\begin{pmatrix} I\\ \left(P_2+\hat{P}\right)\\ I \end{pmatrix} \hat{x}$$

where

$$\hat{L} = \begin{pmatrix} I & \beta B (J'J)^{-1} B' \\ 0 & \beta \left[A' - H'J (J'J)^{-1} B' \right] \end{pmatrix}$$
$$\hat{N} = \begin{pmatrix} A - B (J'J)^{-1} J'H + CK^* & 0 \\ - \left[H'H - H'J (J'J)^{-1} JH \right] & I \end{pmatrix},$$

and (b) to show that the implied law of motion for \hat{x}_{t+1} agrees with

$$\hat{x}_{t+1} = A^* \hat{x}_t. \tag{6.6.7}$$

In constructing the deflating subspace in part (a), we will show that

$$P_2 + \hat{P} = P.$$

This can be done by imitating the argument that $\{w_{t+1}(F^*)\} = \{w_{t+1}^*\}$ but reversing the roles of $\{w_{t+1}\}$ and $\{u_t\}$. So we impose $w_{t+1} = K^*x_t$. It follows that P can indeed be used to represent the $\sqrt{\beta}$ deflating subspace and that the implied evolution for $\{\hat{x}_{t+1}\}$ is given by (6.6.7) as required by part (b).

Thus we have shown that the $\sqrt{\beta}$ deflating subspace can indeed be uncoupled. By initializing $\hat{x}_0 = x_0$, it follows that $\hat{x}_t = x_t$. Moreover, the optimized objective coincides with $C(F^*, x_0)$. Thus

$$C(F, x_0) \le \hat{C}(F, x_0) \le C(F^*, x_0).$$

Theorem 6.6.1 imposes different assumptions from those in Theorem 6.5.1 in order to reflect the change in the minimizing player's view about the maximizing player when we move from game 1 (SEQ) to the Stackelberg game 2 (STACK). Nevertheless, formulas (6.5.9) for the equilibrium objects (K^*, F^*) of game SEQ describe the equilibrium of the Stackelberg multiplier game, and the same notion of stability prevails.

6.7. Game 3: Markov perfect multiplier game

Relative to game 1, game 2 increased the minimizing player's power over the maximizing player by confronting the minimizing player with the law of motion

$$x_{t+1} = (A - BF^*) x_t + Cw_{t+1}$$

$$u_t = -F^* x_t.$$
(6.7.1)

This tells the minimizing player that he can influence future u_s 's by choice of w_{s+1} . Meanwhile in game 2, the maximizing player still faces a sequence w that is exogenous to its choices.

Game 3 continues to confront the minimizing player with (6.7.1) as in game 2, but enhances the power of the maximizing player by confronting him with

$$x_{t+1} = (A + CK^*) x_t + Bu_t$$

$$w_{t+1} = K^* x_t.$$
(6.7.2)

This lets the maximizing player influence future w_s 's. This timing protocol leads to what we refer to as a *Markov perfect multiplier game* (MARKOV).¹¹

 $^{^{11}}$ As we shall see, this game connects directly to the discounted risk-sensitivity criterion of Hansen and Sargent (1995) that is described in chapter 7 on page 195.

Although we have augmented the maximizing player's power, the equilibrium outcome of our game 3 matches that of the Stackelberg multiplier game 2.

Definition 6.7.1. An equilibrium of the Markov perfect multiplier game (MARKOV) is a pair of strategies $u_t = -F^*x_t$, $w_{t+1} = K^*x_t$ such that

(a) Given K^* , $u_t = -F^*x_t$ maximizes (6.5.1), subject to

$$x_{t+1} = Ax_t + Bu_t + Cw_{t+1}.$$
(6.7.3)

(b) Given F^* , $w_{t+1} = -K^* x_t$ minimizes (6.5.1) subject to (6.7.3).

Associated with a Markov perfect multiplier game is the following pair of Bellman equations

$$-x'P^*x = \max_{u} \left[-(Hx + Ju)'(Hx + Ju) + \beta\theta w^{*'}w^* - \beta y'P^*y \right] (6.7.4a)$$

$$y = (A + CK^{*}) x + Bu$$
(6.7.4b)

* K^{*}

$$w^* = K^* x \tag{6.7.4c}$$

$$-x'P^*x = \min_{w} \left[-(Hx + Ju^*)'(Hx + Ju^*) + \beta\theta w'w - \beta y'P^*y \right] (6.7.5a)$$

$$y = (A - BF^*)x + Cw$$
(6.7.5b)

$$u^* = -F^*x, (6.7.5c)$$

The (P^*, K^*, F^*) that form an equilibrium of the MARKOV game also solve the following closely related zero-sum game:

$$-x'P^*x = \max_{u} \min_{w} \left[-(Hx + Ju)'(Hx + Ju) + \beta\theta w'w - \beta y'P^*y \right]$$
(6.7.6)

where the maximization is subject to

$$y = Ax + Bu + Cw.$$

Equilibrium strategies are

$$u = -F^*x,$$

and

$$w = K^* x.$$

Though it yields the same equilibrium strategies and outcome path $x_{t+1} = (A - BF^* + CK^*)x_t$, notice that game (6.7.6) has a slightly different timing protocol from game (6.7.4)–(6.7.5). In (6.7.6), within each period, the *w*-player moves after the *u*-player, while (6.7.4)-(6.7.5) incorporates simultaneous moves within periods.

The (P^*, K^*, F^*) associated with the equilibrium of the Stackelberg multiplier game solves (6.7.4) and (6.7.5), and thereby determines an equilibrium of the MARKOV game. We summarize the connections between an equilibrium of game MARKOV and an equilibrium of game STACK in

Theorem 6.7.1. The (P^*, K^*, F^*) associated with the equilibrium of the Stackelberg multiplier game also describe the equilibrium of the Markov perfect multiplier game (MARKOV).

Proof. The required marginal conditions match.

The functional equation (6.7.6) leads directly to computing the equilibrium by iterating to convergence on Hansen and Sargent's (1995) composite operator $T \circ \mathcal{D}$. The *T* operator represents the maximization over *u* and the \mathcal{D} operator the minimization over *w* in (6.7.6).

6.8. Relation between multiplier and constraint problems

Our three multiplier games have identical outcomes and equilibrium representations. The following two propositions link the multiplier game to a constraint formulation.

Theorem 6.8.1. Suppose that there exists a solution u^*, w^* to the robust multiplier problem in Definition 6.2.2. Then u^* also solves the constraint robust control problem with $\eta = \eta^* = R(w^*)$, where R(w) is defined by (6.2.4).

Theorem 6.8.2. Suppose that u^*, w^* solve the constraint robust control problem in Definition 6.2.1 for $\eta = \eta^*$. Then there exists a θ^* such that the robust multiplier and constraint problems have the same solution.

The propositions follow from the Lagrange multiplier theorem (Luenberger (1969), pp. 216-221). Chapter 7 develops this connection in more detail in

the context of a frequency domain specification of Stackelberg multiplier and constraint games. 12

Theorem 6.8.1 shows how to construct the specification error η associated with a given multiplier θ . In the next section, we use this finding to describe a sense in which the constraint problem is recursive.

6.9. Recursivity of the constraint game

The Bellman equation (6.7.6) indicates directly that the solution of the multiplier problem is time consistent. It requires more of an argument to verify a sense in which the solution of the constraint problem is time consistent because we must identify an additional state variable and describe its law of motion. This section describes a recursive formulation of the constraint problem stated in Definition 6.2.1.¹³ We define a time t version of continuation entropy (6.2.4) as an additional state variable:

$$R_t(w) = \sum_{\tau=1}^{\infty} \beta^{\tau} w'_{t+\tau} w_{t+\tau}.$$

Evidently, $R_t(w)$ satisfies the recursion

$$R_{t}(w) = \beta w_{t+1}' w_{t+1} + \beta R_{t+1}(w) \, .$$

Let $V(x,\eta)$ be the value function for the constraint problem (6.2.5) starting from initial state $x = x_0$ and initial value of entropy η . For the constraint problem, the counterpart to Bellman equation (6.7.6) is¹⁴

$$V(x,\eta) = \sup_{u} \inf_{w,\check{\eta}} \left[-z'z + \beta V(\check{x},\check{\eta}) \right]$$
(6.9.2)

$$V(x,\eta) = \sup_{u} \inf_{w,\breve{\eta}(\epsilon)} \left[-z'z + \beta EV(\check{x},\breve{\eta}) \right]$$
(6.9.1)

where the extremization is subject to

$$\dot{x} = Ax + Bu + C (\epsilon + w)$$

$$\eta = \beta w' w + \beta \check{\eta} (\epsilon),$$

¹² Also see Hansen, Sargent, Turmuhambetova, and Williams (2001).

 $^{1^3}$ See Hansen, Sargent, Turmuhambetova, and Williams (2001) for an extended discussion of the subject of this section.

¹⁴ If a random vector ϵ is present in the transition law, the Bellman equation becomes

where $(\tilde{\cdot})$ denotes next period's value and the extremization is subject to

$$z = Hx + Ju$$

$$\check{x} = Ax + Bu + Cw$$

$$\eta = \beta w'w + \beta \check{\eta}.$$
(6.9.3)

The last equation of (6.9.3) is a 'promise keeping' constraint on the allocation of entropy η between today's model distortion w'w and the distortion from tomorrow on $\check{\eta}$. The minimizing agent can allocate η over time, but must respect constraint (6.9.3).

The first order necessary condition with respect to $\check{\eta}$ and the envelope condition for η imply that

$$V_{\eta}(x,\eta) = V_{\eta}(\check{x},\check{\eta}). \tag{6.9.4}$$

Further, $V_{\eta}(x, \eta)$ equals minus θ , interpreted as the Lagrange multiplier on the last constraint in (6.9.3). Equation (6.9.4) implies that there is a time-invariant relationship between x and η , which in turn implies that the extremizing choices (u, w) for the right side of (6.9.3) can be expressed as functions of x alone. These equal the functions $u = -F^*x$ and $w = K^*x$ that we computed for the multiplier games for θ being set equal to $V_{\eta}(x, \eta)$.

6.10. Summary and concluding remarks

A robust decision maker fears that his approximating model is misspecified and assumes that misspecification takes the form of nonzero shocks $\{w_{t+1}\}$. To attain a robust decision rule, the decision maker modifies the usual Bellman equation by adding another player ('nature') who, by choosing a nearby model to hurt the decision maker, assists the decision maker to find a robust decision rule. Thus, the decision maker devises a robust decision rule by finding a value function v(x) that solves:

$$v(x) = \max_{u} \min_{w} \left\{ -(Hx + Ju)'(Hx + Ju) + \beta \theta w'w + \beta v(y) \right\}$$
(6.10.1)

where E is the mathematical expectation with respect to the distribution of ϵ and continuation entropy $\check{\eta}(\epsilon)$ is now a function of ϵ .

where

$$y = A_o x + Bu + Cw \tag{6.10.2}$$

and θ satisfies $0 < \underline{\theta} < \theta < \infty$, where a formula for the 'breakdown point' $\underline{\theta}$ is reported on page 184. When $\theta = \infty$, the decision maker has no concern about model misspecification and we are back with the ordinary control problem. When $\theta < +\infty$, the decision maker wants a robust rule. The optimum value function is $v(x) = -x'P^*x$ and is attained by a pair of decision rules $u = -F^*x, w_{t+1} = K^*x$, where P^* solves an adjusted Riccati equation $P^* = T \circ \mathcal{D}(P^*)$ and F^* and K^* depend on P^* .

The robust rule F^* is as easy to compute as a rule without robustness because the Bellman equation (6.10.1) is so simple.

The Bellman equation (6.10.1) embeds the timing protocol of the Markov perfect game, one of three two-player zero-sum games with identical payoffs but differing timing protocols. A remarkable outcome is that all three games have identical equilibrium outcomes and identical recursive equilibrium representations $u_t = -F^*x_t$, $w_{t+1} = K^*x_t$, and $x_{t+1} = (A - BF^* + CK^*)x_t$. The zero-sum feature of the games is the essential element in giving all three games the same equilibrium outcomes and representations. The different games justify alternative algorithms for computing F^* and K^* .

Subsequent chapters make ample use of the identity of outcomes from the three multiplier games with different timing protocols. For example, chapter 7 will focus exclusively on a Stackelberg multiplier game in order to deduce some frequency domain characterizations of robust decision rules. That those frequency domain characterizations apply to our other timing protocols rests on the results of this chapter. By appealing to a version of the Stackelberg problem, chapter 9 characterizes the distortion in the endowment process that allows a permanent income consumer to attain a robust decision rule by engaging in a form of precautionary saving. As another example, chapter 16 will formulate a Ramsey problem by using the freedom that the results in this chapter give us to use stacked system of Euler equations formed from the first-order conditions of game 1 (SEQ) to characterize the best response of the robust decision makers who compose the competitive private sector. The validity of that approach depends on the identity of outcomes across our three timing protocols.

Chapter 7 analyzes a multiplier game in the frequency domain under a Stackelberg timing protocol. Working in the frequency domain is equivalent with working in a space of sequences, so in chapter 7 we are compelled to adopt the Stackelberg formulation.

A. Certainty equivalence

A certainty equivalence result that we shall use extensively, especially in chapter 9, has a very similar structure to Theorem 6.5.2. A wide class of decision problems in macroeconomics automatically take the form of a discounted augmented linear regulator where the objective function is

$$-\begin{bmatrix}x_{1t}\\x_{2t}\end{bmatrix}'\begin{bmatrix}P_{11}&P_{12}\\P_{21}&P_{22}\end{bmatrix}\begin{bmatrix}x_{1t}\\x_{2t}\end{bmatrix}-\rho=E\sum_{t=0}^{\infty}\beta^{t}\left\{-\begin{bmatrix}x_{1t}\\x_{2t}\end{bmatrix}'\begin{bmatrix}R_{11}&R_{12}\\R_{21}&R_{22}\end{bmatrix}\begin{bmatrix}x_{1t}\\x_{2t}\end{bmatrix}-u_{t}Qu_{t}\right\}$$

$$(6.A.1)$$

and the transition law is

$$\begin{bmatrix} x_{1t+1} \\ x_{2t+1} \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{bmatrix} \begin{bmatrix} x_{1t} \\ x_{2t} \end{bmatrix} + \begin{bmatrix} 0 \\ C_2 \end{bmatrix} \epsilon_{t+1}$$
(6.A.2)

where ϵ_{t+1} is an i.i.d. random vector with mean zero and identity covariance matrix. The optimal (non robust) decision rule is

$$u_t = -F_1 x_{1t} - F_2 x_{2t} \tag{6.A.3}$$

where F_1 can and F_2 can be computed recursively as in the augmented linear regulator in chapter 3; F_1 is the feedback part and F_2 is the feedforward part.

For a given $\theta \in (\underline{\theta}, \infty)$, we can solve a robust linear regulator and obtain another decision rule

$$u_t = -\tilde{F}_1 x_{1t} - \tilde{F}_2 x_{2t} \tag{6.A.4}$$

of the form (6.A.3) where now \tilde{F}_1 and \tilde{F}_2 depend on θ and C_2 . Let $w_{t+1} = \begin{bmatrix} \tilde{K}_1 & \tilde{K}_2 \end{bmatrix} \begin{bmatrix} x_{1t} \\ x_{2t} \end{bmatrix}$ be the associated worst case shock. We can apply the method used in Theorem 6.5.2 to construct a law of motion that is distorted relative to the approximating model (6.A.2) and for which an ordinary (non robust) decision rule matches the robust rule (6.A.4) for a given θ . Form the law of motion for the synthetic variable

$$\begin{bmatrix} \hat{x}_{1t+1} \\ \hat{x}_{2t+1} \end{bmatrix} = \begin{bmatrix} A_{11} - B_1 F_1 & A_{12} - B_1 F_2 \\ C_2 K_1 & A_{22} + C_2 K_2 \end{bmatrix} \begin{bmatrix} \hat{x}_{1t} \\ \hat{x}_{2t} \end{bmatrix}$$
(6.A.5)

Now alter the law of motion for x_1 in problem (6.A.1)-(6.A.2) to be

$$x_{1t+1} = A_{11}x_{1t} + A_{12}\hat{x}_{2t} + B_1u_t \tag{6.A.6}$$

and use \hat{x}_{2t} to replace x_{2t} in the objective (6.A.1). Solve the ordinary control problem with (6.A.5), (6.A.6) replacing (6.A.2). This again is a discounted augmented linear regulator problem. The decision rule is

$$u_t = -F_1 x_{1t} - \hat{F}_{21} \hat{x}_{1t} - \hat{F}_{22} \hat{x}_{2t}.$$
(6.A.7)

Equating \hat{x}_t to \hat{x}_t , we obtain

$$u_t = -\left(F_1 + \hat{F}_{21}\right) x_{1t} - \hat{F}_{22} x_{2t}.$$
(6.A.8)

Then

$$\tilde{F}_1 = F_1 + \hat{F}_{21}, \tilde{F}_2 = \hat{F}_{22}.$$
 (6.A.9)

Notice the presence of role of the feedback component F_1 from the ordinary control problem. Robustness appears because of the distortion of the law of motion for the x_{2t} component, which enters through the \hat{F}_{2i} , i = 1, 2 terms.

B. Useful formulas

This appendix uses game (6.10.1) to provide two sets of convenient formulas for computing the robust decision rule. For the purpose of displaying these formulas, notice that the one-period loss function in (6.10.1) can be represented as

$$\begin{split} r\left(x,u\right) &\equiv \left(Hx+Ju\right)'\left(Hx+Ju\right) \\ &= \begin{bmatrix} x\\ u \end{bmatrix}' \begin{bmatrix} H'H & H'J\\ J'H & J'J \end{bmatrix} \begin{bmatrix} x\\ u \end{bmatrix} \\ &\equiv \begin{bmatrix} x\\ u \end{bmatrix}' \begin{bmatrix} \overline{Q} & W\\ W' & R \end{bmatrix} \begin{bmatrix} x\\ u \end{bmatrix}, \end{split}$$

where $\overline{Q} = H'h, W = H'J, R = J'J$. As in chapter 3, we transform the problem to one that eliminates cross-products between states and controls. Define

$$Q = \overline{Q} - WR^{-1}W'$$

$$\tilde{A} = A - BR^{-1}W'$$

$$\tilde{u}_t = u_t + R^{-1}W'x_t.$$

(6.B.1)

Then

$$x_{t+1} = \tilde{A}x_t + B\tilde{u}_t + Cw_{t+1}$$
(6.B.2)

and

$$\tilde{r}(x,\tilde{u}) = r(x,u) = \begin{bmatrix} x\\ \tilde{u} \end{bmatrix}' \begin{bmatrix} Q & 0\\ 0 & R \end{bmatrix} \begin{bmatrix} x\\ \tilde{u} \end{bmatrix}.$$
(6.B.3)
The Bellman equation (6.10.1) is equivalent with

$$-x'Px = \max_{\tilde{u}} \min_{w} \left\{ -\tilde{r} \left(x, \tilde{u} \right) + \beta \theta w'w - \beta y'Py \right\}$$
(6.B.4)

where

$$y = Ax + B\tilde{u} + Cw. \tag{6.B.5}$$

In the problem on the right of (6.B.4), the minimizing agent moves second, taking as given the feedback rule $\tilde{u} = -Fx$ chosen by the maximizing agent. By working backwards, we break the problem on the right of (6.B.4) into these two parts:

1. The problem for the minimizing agent reduces to

$$\mathcal{J} = \min_{w} \left[\theta w' w - y' P y \right] \tag{6.B.6}$$

subject to

$$y = \check{A}x + Cw \tag{6.B.7}$$

where $\check{A} = \tilde{A} - BF$ and F is to be chosen in the problem in part 2. The minimizing w is

$$w = \theta^{-1} \left(I - \theta^{-1} C' P C \right)^{-1} C' P \check{A} x.$$
 (6.B.8)

Let

$$\mathcal{D}(P) = P + PC \left(\theta I - C'PC\right)^{-1} C'P.$$
(6.B.9)

The minimized value of the problem can be expressed as

$$\mathcal{J} = -x'\tilde{A}'\mathcal{D}(P)\tilde{A}x$$
$$\mathcal{I} = -u'\mathcal{D}(P)u$$

or as

$$\mathcal{J} = -y'\mathcal{D}\left(P\right)y\tag{6.B.10}$$

where in (6.B.10), y is to be evaluated under the approximating model y = Ax, not under the distorted model (6.B.7). Under the approximating model, (6.B.10) is a conservative continuation value for the problem of the maximizing agent.¹⁵

2. Part 2 of the problem hands this conservative valuation function and the approximating model to the maximizing agent. Working backwards, the problem of the maximizing agent can be expressed as

$$\max_{\tilde{u}} \left[-x'Qx - \tilde{u}'R\tilde{u} - \beta y'\mathcal{D}\left(P\right)y \right]$$
(6.B.11)

subject to

$$y = \tilde{A}x + B\tilde{u}. \tag{6.B.12}$$

¹⁵ In chapter 7, the operator \mathcal{D} is used again to characterize risk-sensitive preferences. See page 195.

Notice that (6.B.12) is the approximating model and that allowance for distortions occurs only through the presence of $\mathcal{D}(P)$ on the right side of (6.B.11). The solution to this problem is found by taking one step on the usual Riccati equation, with $\mathcal{D}(P)$ as the terminal value function. Thus, define the operators

$$\mathcal{F}(\Omega) = \beta \left[R + \beta B' \Omega B \right]^{-1} B' \Omega \tilde{A} \tag{6.B.13}$$

$$T(P) = Q + \beta \tilde{A}' \left(P - \beta P B \left(R + \beta B' P B \right)^{-1} B' P \right) \tilde{A}.$$
(6.B.14)

Substituting in the definitions of Q and R, T can also be expressed as

$$T(P) = H'H - H'J(J'J)^{-1}J'H + \beta \left(A - B(J'J)^{-1}J'H\right)' \times \left(P - \beta PB(J'J + \beta B'PB)^{-1}B'P\right) \left(A - B(J'J)^{-1}J'H\right).$$
(6.B.15)

(6.B.15) Then the solution of problem (6.B.11) is $\tilde{u} = -Fx$ where $F = \mathcal{F} \circ \mathcal{D}(P)$. The maximizing value of (6.B.11) is $-x'T \circ \mathcal{D}(P)x$. Notice that $u_t = \tilde{u}_t - R^{-1}Wx_t = -(F + (J'J)^{-1}J'H)x_t$.

We can iterate on these two sub problems to find the solution to (6.B.4).¹⁶ Let P be the fixed point of iterations on $T \circ \mathcal{D}$:

$$P = T \circ \mathcal{D}(\mathcal{P}). \tag{6.B.16}$$

Then the solution of (6.B.4), (6.B.5) is

$$\tilde{u} = -Fx \tag{6.B.17}$$

$$w = Kx, \tag{6.B.18}$$

where

$$F = \mathcal{F} \circ \mathcal{D} \left(P \right) \tag{6.B.19}$$

$$K = \theta^{-1} \left(I - \theta^{-1} C' P C \right)^{-1} C' P \left[\tilde{A} - BF \right].$$

$$(6.B.20)$$

Here T is the usual operator associated with taking one-step on the Bellman equation without a preference for robustness; it represents optimization with respect to u. The operator \mathcal{D} reflects minimization with respect to w. When $\theta = +\infty$, $\mathcal{D}(P) = P$, and we get the usual optimal rule for a linear-quadratic dynamic program. When $\underline{\theta} \leq \theta < \infty$, we get a robust decision rule, where $\underline{\theta}$ is a lower bound on admissible parameters θ . We shall give a formula for $\underline{\theta}$ in formula (7.3.13) on page 184.

¹⁶ In chapter 7 we show how the two operators are related to the discounted risk-sensitivity criterion of Hansen and Sargent (1995).

6.B.1. A single Riccati equation

A robust decision rule can also be computed simply by solving an optimal linear regulator problem.¹⁷ This can be established in the following way.

By writing iterations $P_{k+1} = T \circ \mathcal{D}(P_k)$ and rearranging, the matrix P in the value function -x'Px can be expressed as the fixed point of iterations on the Riccati equation¹⁸

$$P_{k+1} = \tilde{A}' \left(\left(\beta P_k \right)^{-1} + BR^{-1}B' - \theta^{-1}\beta^{-1}CC' \right) \tilde{A} + Q.$$
 (6.B.21)

This equation can also be represented as

$$P_{k+1} = Q + \tilde{A}^{*\prime} \left(P_k^{-1} + \tilde{J} \right)^{-1} \tilde{A}^*, \qquad (6.B.22)$$

where $\tilde{J} = B^* R^{-1} B^{*'} - \theta^{-1} CC', B^* = \beta^{.5} B, \tilde{A}^* = \beta^{.5} \tilde{A}$. Equation (6.B.22) is in a form to which the *doubling algorithm* described in chapter 3 applies.¹⁹ Notice that (6.B.21) is the Riccati equation associated with an ordinary optimal linear regulator problem with controls $\begin{bmatrix} u \\ w \end{bmatrix}$ and penalty matrix on those controls appearing in the criterion function of $\begin{bmatrix} R & 0 \\ 0 & -\beta\theta I \end{bmatrix}$. Therefore, the robust rules for u_t and the associated worst case shock can be computed directly from the associated ordinary linear regulator problem. It can be checked that the right side of (6.B.21) implements one step on $T \circ \mathcal{D}$. The Riccati equation (6.B.21) is the one associated with the modified linear regulator used in chapter 2 on page 37 to compute a robust rule and the worst case shock.

6.B.2. Robustness bound

The inner problem (6.B.6) inspires a robustness bound for continuation values. Thus, (6.B.6) implies

$$-x'A'\mathcal{D}(P)Ax = \min_{w} \left[\theta w'w - y'Py\right] \le \theta w'w - y'Py \qquad (6.B.23)$$

where y is evaluated under the distorted model y = Ax + Cw. Inequalities (6.B.23) imply

$$-y'Py \ge -x'A'\mathcal{D}(P)Ax - \theta w'w. \tag{6.B.24}$$

The left side is evaluated under a distorted model y = Ax + Cw while the quadratic form in x on the right is a conservative estimate of the continuation value of the state

¹⁷ See chapter 2, page 37.

¹⁸ See Hansen and Sargent (2004, Recursive models).

 $^{^{19}}$ The Matlab program doublex9.m computes the solution using the doubling algorithm.

y under the approximating model y = Ax.²⁰ Inequality (6.B.24) says that the continuation value is at least as great as a conservative estimate of the continuation value under the approximating (w = 0) model, minus θ times the measure of model misspecification w'w. The parameter θ influences the conservative-adjustment operator \mathcal{D} and also determines the rate at which the bound deteriorates with misspecification. Lowering θ lowers the rate at which the bound deteriorates with misspecification. Thus, (6.B.24) provides a sense in which lower values of θ provide more conservative and also more robust estimates of continuation utility.

6.B.3. Special case: a pure forecasting problem

Here is an example of a pure forecasting problem in which the absence of a control eliminates the maximization part of (6.B.4). The following state space system governs consumption and bliss consumption:

$$x_{t+1} = Ax_t + Cw_{t+1}$$

$$c_t = H_c x_t$$

$$b_t = H_b x_t$$
(6.B.25)

where c_t is an exogenous scalar consumption process, b_t is a bliss level of consumption, and w_{t+1} is a specification error sequence. To attain a conservative way of evaluating $-\sum_{t=0}^{\infty} \beta^t (c_t - b_t)^2$, we compute

$$-x_0' P x_0 = \min_{\{w_{t+1}\}} -\sum_{t=0}^{\infty} \beta^t \left[x_t' H' H x_t - \beta \theta w_{t+1}' w_{t+1} \right]$$
(6.B.26)

subject to (6.B.25), where $H = H_c - H_b$. For this special case, the absence of a control causes the operator T defined in (6.B.14) to simplify to

$$T(P) = H'H + \beta A'PA. \tag{6.B.27}$$

The matrix P in (6.B.26) is the fixed point of iterations on $T \circ \mathcal{D}$. The minimizer of (6.B.26) is given by (6.B.8), or w = Kx, where K is defined implicitly by (6.B.8). It follows from our earlier characterizations of K and $P = T \circ \mathcal{D}(P)$ that

$$-x_0'Px_0 = -\sum_{t=0}^{\infty} \beta^t x_t' H' Hx_t$$

where the right side is computed using the distorted law of motion

$$x_{t+1} = (A + KC) x_t.$$

²⁰ That is, when w = 0, $-y' \mathcal{D}(P)y$ understates the continuation value.

C. Details of a proof of Theorem 6.5.2

This appendix verifies a key assertion made in the proof of Theorem 6.5.2.²¹ To verify that $\hat{P} = P^* - \bar{P}$, we make use of the following identities that characterize \hat{P}, P^* , and \bar{P} :

1. The matrix \hat{P} solves:

$$\beta \left(A' - H'J \left(J'J \right)^{-1} B' \right) \left(\bar{P} - \beta \bar{P}B \left(J'J + \beta B' \bar{P}B \right)^{-1} B' \bar{P} \right) \hat{A} + \beta \left(A' - H'J \left(J'J \right)^{-1} B' \right) \left(\hat{P} - \beta \bar{P}B \left(J'J + \beta B' \bar{P}B \right)^{-1} B' \hat{P} \right) A^* = \hat{P}$$

where:

$$A^* = \left(I + \left[\beta B \left(J'J\right)^{-1} B' - \frac{1}{\theta} C C'\right] P^*\right)^{-1} \left[A - B \left(J'J\right)^{-1} J' H\right]$$
$$\hat{A} = \frac{1}{\theta} C C' P^* A^*$$

Therefore, the matrix \hat{P} solves:

$$\beta \left(A' - H'J \left(J'J \right)^{-1} B' \right) \left[\bar{P} \frac{1}{\theta} CC' P^* - \beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} B' \bar{P} \frac{1}{\theta} CC' P^* + \hat{P} - \beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} B' \hat{P} \right] A^* = \hat{P}$$

which yields:

$$\beta \left(A' - H'J \left(J'J \right)^{-1} B' \right) \\ \left[\bar{P} \frac{1}{\theta} CC' P^* - \beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} B' \bar{P} \frac{1}{\theta} CC' P^* + \hat{P} - \beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} B' \hat{P} \right] \\ \times \left(I + \left[\beta B \left(J'J \right)^{-1} B' - \frac{1}{\theta} CC' \right] P^* \right)^{-1} \left[A - B \left(J'J \right)^{-1} J' H \right] = \hat{P}$$
(6.C.1)

2. The matrix P^* solves:

$$\beta \left[A' - H'J (J'J)^{-1} B' \right] P^* \left(I + \left[\beta B (J'J)^{-1} B' - \frac{1}{\theta} CC' \right] P^* \right)^{-1} \left[A - B (J'J)^{-1} J' H \right] + H'H - H'J (J'J)^{-1} J'H = P^*$$
(6.C.2)

where:

$$\begin{pmatrix} I + \begin{bmatrix} \beta B (J'J)^{-1} B' - \frac{1}{\theta} CC' \end{bmatrix} P^* \end{pmatrix}^{-1} = \begin{bmatrix} I - \beta (\begin{bmatrix} B & C \end{bmatrix}) \begin{pmatrix} \begin{bmatrix} JJ' + \beta B'P^*B & \beta B'P^*C \\ \beta C'P^*B & -\beta \theta I + \beta C'P^*C \end{bmatrix} \end{pmatrix}^{-1} \begin{pmatrix} \begin{bmatrix} B' \\ C' \end{bmatrix} P^* \end{bmatrix}$$

 $^{^{21}}$ We are very grateful to Tomasz Piskorski for his help in verifying these equalities.

3. The matrix \overline{P} solves:

$$\beta \left(A' - H'J \left(J'J \right)^{-1} B' \right) \left(\bar{P} - \beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} B' \bar{P} \right) \\ \times \left(A - B \left(J'J \right)^{-1} J' H \right) + H' H - H'J \left(J'J \right)^{-1} J' H = \bar{P}.$$
(6.C.3)

We weave together these three facts to compose the following *Proof.* Subtracting (6.C.3) from (6.C.2) yields:

$$\beta \left[A' - H'J \left(J'J \right)^{-1} B' \right] \left(P^* \left(I + \left[\beta B \left(J'J \right)^{-1} B' - \frac{1}{\theta} C C' \right] P^* \right)^{-1} - \left(\bar{P} - \beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} B' \bar{P} \right) \right) \left[A - B \left(J'J \right)^{-1} J' H \right] = P^* - \bar{P}$$

which is equivalent to:

$$\beta \left[A' - H'J \left(J'J \right)^{-1} B' \right]$$

$$\times \left(P^* - \left(\bar{P} - \beta \bar{P}B \left(J'J + \beta B' \bar{P}B \right)^{-1} B' \bar{P} \right) \left(I + \left[\beta B \left(J'J \right)^{-1} B' - \frac{1}{\theta} CC' \right] P^* \right) \right)$$

$$\times \left(I + \left[\beta B \left(J'J \right)^{-1} B' - \frac{1}{\theta} CC' \right] P^* \right)^{-1} \left[A - B \left(J'J \right)^{-1} J' H \right] = P^* - \bar{P}$$
or

or

$$\beta \left[A' - H'J \left(J'J \right)^{-1} B' \right] Y$$

$$\times \left(I + \left[\beta B \left(J'J \right)^{-1} B' - \frac{1}{\theta} CC' \right] P^* \right)^{-1} \left[A - B \left(J'J \right)^{-1} J'H \right] = P^* - \bar{P}$$
(6.C.4)

where:

$$Y = P^{*} - \left(\bar{P} - \beta\bar{P}B\left(J'J + \beta B'\bar{P}B\right)^{-1}B'\bar{P}\right)\left(I + \left[\beta B\left(J'J\right)^{-1}B' - \frac{1}{\theta}CC'\right]P^{*}\right).$$

Note that:

$$Y = P^{*} - \left(\bar{P} - \beta\bar{P}B\left(J'J + \beta B'\bar{P}B\right)^{-1}B'\bar{P}\right)\left(I + \left[\beta B\left(J'J\right)^{-1}B' - \frac{1}{\theta}CC'\right]P^{*}\right)$$

$$= \left(P^{*} - \bar{P}\right) - \bar{P}\beta B\left(J'J\right)^{-1}B'P^{*} + \bar{P}\frac{1}{\theta}CC'P^{*} + \beta\bar{P}B\left(J'J + \beta B'\bar{P}B\right)^{-1}B'\bar{P}$$

$$+ \beta\bar{P}B\left(J'J + \beta B'\bar{P}B\right)^{-1}B'\bar{P}\beta B\left(J'J\right)^{-1}B'P^{*} - \beta\bar{P}B\left(J'J + \beta B'\bar{P}B\right)^{-1}B'\bar{P}\frac{1}{\theta}CC'P^{*}$$

$$= \bar{P}\frac{1}{\theta}CC'P^{*} - \beta\bar{P}B\left(J'J + \beta B'\bar{P}B\right)^{-1}B'\bar{P}B\left(J'J + \beta B'\bar{P}B\right)^{-1}B'\bar{P}\beta B\left(J'J\right)^{-1}B'$$

$$+ \beta\bar{P}B\left(J'J + \beta B'\bar{P}B\right)^{-1}B'\bar{P} + \beta\bar{P}B\left(J'J + \beta B'\bar{P}B\right)^{-1}B'\bar{P}\beta B\left(J'J\right)^{-1}B'$$

$$- \bar{P}\beta B\left(J'J\right)^{-1}B'P^{*}$$

174

$$Y = \bar{P} \frac{1}{\theta} CC' P^* - \beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} B' \bar{P} \frac{1}{\theta} CC' P^* + \left(P^* - \bar{P} \right) + Z \quad (6.C.5)$$
where:

$$Z = \beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} B' \bar{P} + \beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} B' \bar{P} \beta B \left(J'J \right)^{-1} B'$$

$$- \bar{P} \beta B \left(J'J \right)^{-1} B' P^*$$
Now note

$$Z = \beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} B' \bar{P} + \beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} B' \bar{P} \beta B \left(J'J \right)^{-1} B'$$

$$- \bar{P} \beta B \left(J'J \right)^{-1} B' P^* = -\beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} B' \left(P^* - \bar{P} \right)$$

$$+ \beta \bar{P} B \left(J'J \right)^{-1} B' P^* = -\beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} B' \left(P^* - \bar{P} \right)$$

$$+ \left[\beta \bar{P} B \left(J'J \right)^{-1} B' P^* = -\beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} B' \left(P^* - \bar{P} \right)$$

$$+ \left[\beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} \left(I + \beta B' \bar{P} B \left(J'J \right)^{-1} \right) - \beta \bar{P} B \left(J'J \right)^{-1} \right] B' P^*.$$
(6.C.6)

Using the fact that $I + \beta B' \overline{P}B(J'J)^{-1} = (J'J + \beta B' \overline{P}B)(J'J)^{-1}$ gives us that:

$$\left[\beta\bar{P}B\left(J'J+\beta B'\bar{P}B\right)^{-1}\left(I+\beta B'\bar{P}B\left(J'J\right)^{-1}\right)-\beta\bar{P}B\left(J'J\right)^{-1}\right]B'P^{*} = \left[\beta\bar{P}B\left(J'J+\beta B'\bar{P}B\right)^{-1}\left(J'J+\beta B'\bar{P}B\right)\left(J'J\right)^{-1}-\beta\bar{P}B\left(J'J\right)^{-1}\right]B'P^{*} = \left(\beta\bar{P}B\left(J'J\right)^{-1}-\beta\bar{P}B\left(J'J\right)^{-1}\right)P^{*} = 0.$$

$$(6.C.7)$$

Substituting (6.C.7) back to (6.C.6) yields:

$$Z = -\beta \bar{P}B \left(J'J + \beta B'\bar{P}B \right)^{-1} B' \left(P^* - \bar{P} \right).$$
(6.C.8)
s) into (6.C.5) violds:

Substituting (6.C.8) into (6.C.5) yields:

$$Y\bar{P}\frac{1}{\theta}CC'P^{*} - \beta\bar{P}B\left(J'J + \beta B'\bar{P}B\right)^{-1}B'\bar{P}\frac{1}{\theta}CC'P^{*} + (P^{*} - \bar{P}) - \beta\bar{P}B\left(J'J + \beta B'\bar{P}B\right)^{-1}B'(P^{*} - \bar{P}).$$
(6.C.9)

Finally substituting (6.C.9) into (6.C.4) yields:

$$\beta \left[A' - H'J \left(J'J \right)^{-1} B' \right] \\ \times \left[\bar{P} \frac{1}{\theta} CC' P^* - \beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} B' \bar{P} \frac{1}{\theta} CC' P^* + \left(P^* - \bar{P} \right) \right] \\ -\beta \bar{P} B \left(J'J + \beta B' \bar{P} B \right)^{-1} B' \left(P^* - \bar{P} \right) \right] \\ \times \left(I + \left[\beta B \left(J'J \right)^{-1} B' - \frac{1}{\theta} CC' \right] P^* \right)^{-1} \left[A - B \left(J'J \right)^{-1} J' H \right] = P^* - \bar{P}.$$
(6.C.10)

But this is just Riccati equation (6.C.1) with $\hat{P} = P^* - \bar{P}$, therefore $(P^* - \bar{P})$ solves the Riccati equation for \hat{P} , so $\hat{P} = P^* - \bar{P}$.

Chapter 7 Frequency domain games and criteria for robustness

7.1. Three frequency domain criteria indexed by (x_0, η)

In the two-player zero-sum games of chapter 6, the minimizing player helps the maximizing player analyze the fragility of his decision rule $u_t = -Fx_t$. In this chapter, we take the fruitful point of view that an indirect utility function of the minimizing player forms an intertemporal valuation function that when used by the maximizing player produces a robust decision rule. We use frequency domain decompositions to express two such alterations of the maximizing player's objective function. Frequency domain decompositions of variances and other inner products are widely used in the analysis of covariance-stationary time series. Frequency domain decompositions of objective functions like those in this chapter have been used creatively in dynamic macroeconomic models by Whiteman (1985a, 198XXX) and Otrok (2001bXX).

In the interests of evaluating 'steady state' performance of a maximizing player's rule, we require the maximizing player to choose a time-invariant policy rule $u_t = -Fx_t$. However, our frequency domain calculations allow the minimizing player to choose a sequence of shocks $w \in \mathcal{W}$, putting us into the setting of our Stackelberg game of chapter 6.

By minimizing over \mathcal{W} for alternative settings of the initial condition x_0 and the magnitude of the constraint η measuring the size of the specification error, we obtain three different frequency domain criteria that the maximizing agent can use to evaluate alternative time-invariant F's. The first criterion eschews robustness by setting $\eta = 0$, and leads to a discounted version of the so-called H_2 criterion, the maximization of which leads to an algebraic Riccati equation associated with the steady state of a time-invariant, infinite horizon optimal linear regulator problem. Other assumptions about x_0 and η lead to discounted versions of what are known as the H_{∞} criterion and an 'entropy' criterion. Each of these promotes robust decision rules. The entropy criterion has the same multiplier parameter θ that played such a key role in chapter 6.

– 177 –

Indeed, for the same fixed admissible θ , maximizing the entropy criterion leads to the same decision rule F associated with the robust version of the linear regulator that we obtained in chapter 6.

The frequency domain provides an interesting perspective on fear of model misspecification. By analyzing the entropy criterion, we can show that activating a concern about misspecification of the approximating model yields what looks like 'risk-aversion across frequencies' that manifests a preference for smoothness across frequencies. The H_{∞} criterion can be viewed as a limiting version of the entropy criterion approached when the multiplier θ approaches the critical 'breakdown' value $\underline{\theta}$ that we mentioned in chapter 6. The H_{∞} criterion is expressed in terms of the largest eigenvalue across frequencies of that same frequency-domain decomposition of discounted utility and embodies an extreme preference for smoothness across frequencies.

Undiscounted versions of both the H_{∞} and the entropy criteria exist in the control literature. Our analysis of discounting is an innovative part of this chapter. Accommodating discounting requires that, relative to arguments in the control literature, we must pay special attention to initial conditions.¹

Key findings of this chapter are these: (1) the H_2 criterion gives rise to the optimal linear regulator without robustness; (2) for a given $\theta > \underline{\theta}$, the entropy criterion leads to an infinite-horizon time-invariant discounted robust linear regulator with a value function matrix P associated with the limit of iterations on the composite operator $T \circ \mathcal{D}$ described in chapter 6; (3) the break-down value $\underline{\theta}$ equals the squared value of the H_{∞} criterion.

¹ Our derivation of the entropy criterion will also provide a link to the discounted risksensitivity criterion of Hansen and Sargent (1995).

7.2. The Stackelberg game in the time domain

Throughout this chapter we adopt a timing protocol associated with the Stackelberg robust multiplier problem from chapter 6. After recalling this game in the time domain, we shall describe a frequency domain version.

The game requires the maximizing player to choose a time-invariant decision rule of the form $u_t = -Fx_t$. To attain representations that build in $u_t = -Fx_t$, we substitute this decision rule into (6.2.1) to get the closed-loop law of motion for the state:

$$x_{t+1} = A_F x_t + C w_{t+1}, (7.2.1)$$

where

$$A_F = A - BF. \tag{7.2.2}$$

Under $u_t = -Fx_t$, the target becomes

$$z_t = H_F x_t$$

where $H_F = H - JF$. From chapter 6, recall the spaces

$$W = \{w : \sum_{t=1}^{\infty} \beta^t w'_t w_t < +\infty\}$$

$$\mathcal{F} = \{F : A - BF \text{ has eigenvalues with moduli strictly less than} \frac{1}{\sqrt{\beta}}\}.$$

Definition 7.2.1. The Stackelberg robust constraint problem is to find $(F, \{w_t\}_{t=1}^{\infty})$ that attain

$$\max_{F \in \mathcal{F}} \inf_{w \in W} - \sum_{t=0}^{\infty} \beta^t z'_t z_t$$
(7.2.3)

subject to (7.2.1) and

$$\sum_{t=0}^{\infty} \beta^t w_t' w_t \le \eta + w_0' w_0 \tag{7.2.4a}$$

$$x_0 = Cw_0. (7.2.4b)$$

We use $\max_{\mathcal{F}}$ as a shorthand for $\max_{F \in \mathcal{F}}$, and so on. This game is indexed by two parameters (w_0, η) .

Three versions of the Stackelberg robust constraint problem correspond to different settings of η, w_0 :

- 1. The H_2 problem: set $\eta = 0$, with arbitrary w_0 .
- 2. The H_{∞} problem: set $w_0 = 0$, but let $\eta > 0$ be arbitrary.
- 3. The entropy problem: set arbitrary $w_0 \neq 0$ and arbitrary $\eta > 0$.

The first version makes the inf part trivial, turns the game into a standard single-person linear-quadratic optimum problem, and leads to the so-called H_2 criterion in the frequency domain. The second and third versions induce robust decision rules.

To justify an analysis in the frequency domain, (7.2.4b) restricts the initial condition. The solution of the game under this restriction can be represented recursively as a pair of feedback rules $w_{t+1} = Kx_t, u_t = -Fx_t$.²

7.3. Stackelberg game in frequency domain

7.3.1. Fourier transforms

To formulate the game in the frequency domain, define one-sided Fourier transforms:

$$X(\zeta) \equiv \sum_{t=0}^{\infty} x_t \zeta^t,$$

$$W(\zeta) \equiv \sum_{t=0}^{\infty} w_t \zeta^t,$$

$$Z(\zeta) \equiv \sum_{t=0}^{\infty} z_t \zeta^t,$$

(7.3.1)

where ζ is a complex variable. Then (7.2.1) and (7.3.1) imply that $\zeta^{-1}[X(\zeta) - x_0] = A_F X(\zeta) + \zeta^{-1} C [W(\zeta) - w_0]$. Using (7.2.4b) and solving for $X(\zeta)$ gives

² As can be verified by inspecting the formulas for K, F that we derive later, the solution also solves the multiplier games from chapter 6. The time domain representation of the solution of this multiplier game is therefore valid for an arbitrary initial x_0 .

 $X(\zeta) = (I - \zeta A_F)^{-1} CW(\zeta)$, and hence

$$Z\left(\zeta\right) = G_F\left(\zeta\right) W\left(\zeta\right) \tag{7.3.2}$$

where

$$G_F(\zeta) \equiv H_F \left(I - \zeta A_F \right)^{-1} C$$

is the transfer function from shocks to targets.

Applying Parseval's equality to (7.3.2) gives the following representation:

$$\sum_{t=0}^{\infty} \beta^{t} z_{t}' z_{t} = \int_{\Gamma} W\left(\zeta\right)' G_{F}\left(\zeta\right)' G_{F}\left(\zeta\right) W\left(\zeta\right) d\lambda\left(\zeta\right), \qquad (7.3.3)$$

where the operation ' denotes both matrix transposition and complex conjugation, where the measure λ has a density given by

$$d\lambda\left(\zeta\right) \equiv \frac{1}{2\pi i\sqrt{\beta}\zeta}d\zeta,$$

and where the region of integration is the following circle in the complex plane

$$\Gamma \equiv \{\zeta : |\zeta| = \sqrt{\beta}\}.$$

The region Γ can be parameterized conveniently in terms of $\zeta = \sqrt{\beta} \exp(i\omega)$ for ω in the interval $(-\pi, \pi]$. Here the measure λ satisfies

$$d\lambda\left(\zeta\right) = \frac{1}{2\pi}d\omega.$$

Thus the contour integral on the right side of (7.3.3) can be expressed as:

$$\int_{\Gamma} W(\zeta)' G_F(\zeta)' G_F(\zeta) W(\zeta) d\lambda(\zeta)$$

= $\frac{1}{2\pi} \int_{-\pi}^{\pi} W\left(\sqrt{\beta} \exp(i\omega)\right)' \left\{ G_F\left[\sqrt{\beta} \exp(i\omega)\right]' G_F\left[\sqrt{\beta} \exp(i\omega)\right] \right\} W\left(\sqrt{\beta} \exp(i\omega)\right) d\omega.$
(7.3.4)

We use the contour integral on the left of (7.3.4) to simplify notation.

Parseval's equality also implies

$$\sum_{t=0}^{\infty} \beta^{t} w_{t}' w_{t} = \int_{\Gamma} W\left(\zeta\right)' W\left(\zeta\right) d\lambda\left(\zeta\right).$$
(7.3.5)

7.3.2. H_2 criterion

When $\eta = 0$ in (7.2.4*a*), $W(\zeta) = w_0$ and

$$-\sum_{t=0}^{\infty}\beta^{t}z_{t}'z_{t} = -w_{0}'\left[\int_{\Gamma}G_{F}\left(\zeta\right)'G_{F}\left(\zeta\right)d\lambda\left(\zeta\right)\right]w_{0}.$$

For an arbitrary w_0 , the H_2 problem is to maximize this expression by choosing a feedback rule F. The H_2 criterion can be expressed as

$$H_{2} \equiv -\int_{\Gamma} \operatorname{trace} \left[G_{F}\left(\zeta\right)' G_{F}\left(\zeta\right) \right] d\lambda\left(\zeta\right).$$
(7.3.6)

The same F that maximizes H_2 also solves the standard optimal linear regulator problem. Thus, the H_2 criterion gives a frequency domain expression to the preferences embodied in the optimal linear regulator. We turn next to frequency domain criteria that express a concern about model misspecification.

7.3.3. The Stackelberg game in the frequency game

To represent the Stackelberg game in the frequency domain, we define the following two sets of admissible $W(\zeta)$'s:

 $\mathcal{W}^{a} = \{ W(\zeta) : W(\zeta) \text{ is analytic on the interior of } \Gamma \text{ with coefficients} \\ w_{t} \text{ that are vectors of real numbers and } W(0) = w_{0} \}$

$$\mathcal{W} = \{ W(\zeta) \in \mathcal{W}^a : \sum_{t=0}^{\infty} \beta^t w'_t w_t < \infty \}$$

We use (7.3.3) and (7.3.5) to represent the time-domain Stackelberg robust constraint problem of Definition 7.2.1 as:

Three frequency domain games: Find $(F, W(\zeta))$ that attain

$$\max_{\mathcal{F}} \inf_{\mathcal{W}} - \int_{\Gamma} W\left(\zeta\right)' G_F\left(\zeta\right)' G_F\left(\zeta\right) W\left(\zeta\right) d\lambda\left(\zeta\right)$$
(7.3.7)

subject to

$$\int_{\Gamma} W(\zeta)' W(\zeta) \, d\lambda(\zeta) \le \eta + w_0' w_0. \tag{7.3.8}$$

In the frequency domain, our three versions of the Stackelberg multiplier game are:

- 1. H_2 : set $\eta = 0$, with $W(0) = w_0$ arbitrary.
- 2. H_{∞} : set arbitrary $\eta > 0$ but $W(0) = w_0 = 0$.
- 3. Entropy: set $W(0) = w_0 \neq 0$ and arbitrary $\eta > 0$.

We have observed that version 1 leads to the H_2 criterion (7.3.6) and how under it the best feedback rule F is independent of the initial condition $x_0 = Cw_0$. Version 2 has the side condition that W(0) = 0, but otherwise leaves $W(\zeta)$ free. Version 3 requires $W(0) = w_0 \neq 0$ and also restricts $W(\zeta)$ to keep the associated $\{w_t\}$ sequence zero for t < 0.

7.3.4. Version 2: the H_{∞} Criterion

Let $\rho(\zeta)$ denote the eigenvalues of $G_F(\zeta)'G_F(\zeta)$. The following theorem tells how version 2 of the game leads to the H_{∞} criterion defined as:

$$H_{\infty} \equiv -\sup_{\zeta \in \Gamma} \left[\rho\left(\zeta\right) \right]^{1/2}.$$
(7.3.9)

Theorem 7.3.1. For any $F \in \mathcal{F}$,

$$\inf_{\mathcal{W}} - \int_{\Gamma} W\left(\zeta\right)' G_F\left(\zeta\right)' G_F\left(\zeta\right) W\left(\zeta\right) d\lambda\left(\zeta\right) = -H_{\infty}^2 \eta \qquad (7.3.10)$$

where the infinization is subject to (7.3.8).

Proof. Given $G_F(\zeta)$, for each $\zeta = \sqrt{\beta} \exp(i\omega)$ solve the following eigenvalue problem³

$$G_{F}(\zeta)' G_{F}(\zeta) v = \rho(\zeta) v$$

$$\max_{p} p' V p$$

subject to

$$p'p = 1.$$

³ It may be useful to remind the reader of the principal components problem. Let a be an $(n \times 1)$ random vector with covariance matrix V. The first *principal component* of a is a scalar b = p'a where p is an $(n \times 1)$ vector with unit norm (i.e., p'p = 1), for which the variance of b is maximal. Thus, the first principal component solves the problem:

for the largest eigenvalue $\rho(\zeta)$. This problem has a well defined solution with eigenvalue $\rho(\omega)$ for each $\zeta = \sqrt{\beta} \exp(i\omega)$. Then

$$\int_{\Gamma} W(\zeta)' G_F(\zeta)' G_F(\zeta) W(\zeta) d\lambda(\zeta) \leq \int_{\Gamma} \rho(\zeta) W(\zeta)' W(\zeta) d\lambda(\zeta)$$
$$\leq \sup_{\zeta \in \Gamma} \rho(\zeta) \int_{\Gamma} W(\zeta)' W(\zeta) d\lambda(\zeta)$$
$$\leq \sup_{\zeta \in \Gamma} \rho(\zeta) \eta.$$

The bound on the right side is attained by the limit of a sequence of approximating $\{w_t\}$ sequences described in appendix A.

For technical reasons described in appendix A, the infimum in (7.3.10) is not necessarily attained by an analytic function $W \in \mathcal{W}$.

The square of the optimized H_{∞} criterion equals the lower bound on the set of admissible θ 's alluded to in condition (6.2.9) in chapter 6:

$$\underline{\theta} = \left(\inf_{\mathcal{F}} H_{\infty}\left(F\right)\right)^{2}; \qquad (7.3.13)$$

 $\underline{\theta}$ is called the 'breakdown' value of $\theta.^4$

If version 2 has a maximizer F, that F maximizes (7.3.9). We can drop η from the performance criterion (7.3.9) because it becomes a positive scale factor that is independent of the control law F. This feature emerges from our having imposed the initial condition $w_0 = 0$.

$$(V - \lambda I) p = 0,$$
 (7.3.11)

with the value of the variance of p'b evidently from (7.3.11) being

$$p'Vp = \lambda p'p = \lambda. \tag{7.3.12}$$

Putting a Lagrange multiplier λ on the constraint, the first order conditions for this problem are

Thus (7.3.11) and (7.3.12) indicate that p is the eigenvector of V associated with the largest eigenvalue; and that the variance of b equals the largest eigenvalue λ .

⁴ Whittle (1990) calls $\underline{\theta}$ the point of 'utter psychotic despair'.

7.4. A multiplier game for version 3

The H_2 criterion emerged from ignoring concern about model misspecification by setting $\eta = 0$. Under discounting, the H_{∞} control problem came from allowing model misspecification while setting w_0 to zero. We now consider an intermediate case that allows misspecification but also constrains the malevolent agent to respect the initial condition w_0 . This intermediate case will lead to the following 'entropy criterion' for the maximizing agent:

$$\int_{\Gamma} \log \det \left(\theta I - G'_F G_F \right) d\lambda \left(\zeta \right).$$

We shall explain the appellation 'entropy' in section 7.7. To analyze this case, we formulate the multiplier version of the Stackelberg game in the frequency domain. We let θ be a Lagrange multiplier on the constraint and obtain:

Definition 7.4.1. Lagrangian formulation of Stackelberg constraint game: Find $(\theta, F, W(\zeta))$ that attain

$$\sup_{\theta} \sup_{\mathcal{F}} \inf_{\mathcal{W}} \left[\int_{\Gamma} W' \left(\theta I - G'_F G_F \right) W d\lambda - \theta \left(\eta + w'_0 w_0 \right) \right].$$
(7.4.1)

Here $\eta > 0$ and $w_0 \neq 0.5$

In appendix C, we establish the following things about (7.4.1).

i. Let θ^* be the optimal multiplier for (7.4.1). It satisfies:

$$\theta^* \ge \underline{\theta}.\tag{7.4.2}$$

If (7.4.2) does not hold, the inner $\inf_{W \in \mathcal{W}}$ in (7.4.1) is $-\infty$ independently of the control law F.

ii. When the optimal multiplier θ^* satisfies $\theta^* > \underline{\theta}$, we are led to study the inner two-player zero-sum *Stackelberg multiplier game*:

$$\sup_{\mathcal{F}} \inf_{\mathcal{W}} \int_{\Gamma} W' \left(\theta^* I - G'_F G_F \right) W d\lambda$$
(7.4.3)

This game connects to the single agent decision problem

$$\sup_{\mathcal{F}} \int_{\Gamma} \log \det \left[\theta^* I - G_F(\zeta)' G_F(\zeta) \right] d\lambda(\zeta) , \qquad (7.4.4)$$

⁵ We have already studied the $\eta = 0$ (H₂) and $w_0 = 0$ (H_{∞}) cases.

because the F^* that attains (7.4.4) is the F component of the solution of the two-player multiplier game (7.4.3).

7.5. The multiplier problem and the entropy criterion

To study the $\inf_{\mathcal{W}}$ part of game (7.4.3), we take θ^* , F, and therefore G_F as given.⁶ We refer to the resulting optimization problem as the *multiplier problem* and state it as:

The multiplier problem:

$$\inf_{W,W(0)=w_0} \int_{\Gamma} W\left(\zeta\right)' \left[\theta^* I - G_F\left(\zeta\right)' G_F\left(\zeta\right)\right] W\left(\zeta\right) d\lambda\left(\zeta\right).$$
(7.5.1)

For this problem to have an optimized value that exceeds $-\infty$, we require that $\theta^* I - G'_F G_F$ be positive semidefinite. As a consequence,

$$\theta^* \ge \left[H_\infty\left(F\right)\right]^2,$$

which is a sharper restriction than⁷

$$\theta^* \geq \underline{\theta} = \left[\inf_{\mathcal{F}} H_{\infty}(F)\right]^2.$$

In what follows we strengthen the restriction that $\theta^*I - G'_FG_F$ be positive semidefinite by requiring entropy to be finite:

$$\int_{\Gamma} \log \det \left(\theta^* I - G'_F G_F \right) d\lambda \left(\zeta \right) > -\infty.$$
(7.5.2)

It is necessary to check this condition only at

$$\theta^* = \left[H_\infty\left(F\right)\right]^2,$$

because for larger values of θ^* , (7.5.2) is satisfied automatically. We shall show that for any value of θ^* that exceeds the threshold $\underline{\theta} = [H_{\infty}(F)]^2$, the entropy

 $[\]begin{array}{l} \hline 6 \\ 7 \\ \hline 7 \\ \end{array} \text{Recall that } G_F \equiv H_F (I - \zeta (A - BF))^{-1} C. \end{array}$

measure on the left side of (7.5.2) is closely related to the minimized value of the multiplier problem.

If condition (7.5.2) holds, we can associate choices of θ^* with restrictions on the specification errors. That is, consider the following *constrained worst case* minimization problem:

Constrained worst case problem:

$$\min_{\mathcal{W}} - \int_{\Gamma} W' G'_F G_F W d\lambda$$

subject to

$$\int_{\Gamma} W' W d\lambda \le w'_0 w_0 + \eta.$$

Theorem 7.5.1. For any $\theta^* > [H_{\infty}(F)]^2$, there exists an η such that the multiplier problem and the constrained worst case problem have the same solution.

Proof. See Appendix C.

If the infimum of the multiplier problem is attained for $\theta^* = [H_{\infty}(F)]^2$, then there is a finite η such that the two problems continue to have the same solution. If the infimum is not attained, then any finite η is associated with a multiplier θ^* that exceeds $[H_{\infty}(F)]^2$. Thus we can think of the θ^* 's in the multiplier problem as measuring the size of allowable specification errors.

7.5.1. A robustness bound

For a given decision rule F, the multiplier problem yields an inequality that bounds the rate at which the criterion function deteriorates as specification errors increase. Let \mathcal{J} denote the minimized value of the objective (7.5.1) for the multiplier problem. Then

$$-\int_{\Gamma} W'G'_F G_F W d\lambda \ge \mathcal{J} - \theta^* \int_{\Gamma} W' W d\lambda.$$
(7.5.3)

Inequality (7.5.3) shows that in the absence of specification errors, \mathcal{J} understates the performance of the policy. It also shows how θ^* sets the rate at which the objective function $-\int_{\Gamma} W'G'_F G_F W d\lambda$ deteriorates with model misspecification as measured by $\int W'W d\lambda$. Note how *lowering* θ^* gives more robustness in the sense of less sensitivity of the objective function to misspecifications W.

In the remainder of this section, we study the existence of a solution to the multiplier problem and its relation to the entropy criterion. We return to the Stackelberg multiplier game in the following section.

7.5.2. Entropy is the indirect utility function of the multiplier problem

For establishing our next result, it is convenient to rewrite the multiplier problem as

$$\inf_{W(\zeta)\in\mathcal{W}} \int_{\Gamma} W\left(\zeta\right)' \left[\theta^* I - G_F\left(\zeta\right)' G_F\left(\zeta\right)\right] W\left(\zeta\right) d\lambda\left(\zeta\right)$$
(7.5.4)

subject to

$$\int_{\Gamma} W\left(\zeta\right) d\lambda\left(\zeta\right) = w_0 \neq 0, \tag{7.5.5}$$

and

188

$$\int_{\Gamma} W\left(\zeta\right) \zeta^{j} d\lambda\left(\zeta\right) = 0, \qquad (7.5.6)$$

for j = 1, 2, ... Constraint (7.5.5) can be restated as $W(0) = w_0$. Constraint (7.5.6) states that $w_j = 0$ for j < 0. From the definition of \mathcal{W} , the infimum in (7.5.4) is over $W(\zeta)$ that have coefficients such that $\sum_{t=-\infty}^{\infty} \beta^t w'_t w_t < \infty$.

Theorem 7.5.2. Assume that F and θ^* are such that $\int_{\Gamma} \log \det(\theta^* I - G'_F G_F) d\lambda > -\infty$. Then multiplier problem (7.5.1) has an optimized value function $w'_0 D(0)' D(0) w_0$, where D(0) is nonsingular and independent of w_0 . The minimized value is attained if $\theta^* I - G'_F G_F$ is nonsingular on Γ .

Proof. The solution to the multiplier problem can be found using techniques from linear prediction theory.⁸ We must factor a spectral-density-like matrix:

$$\left[\theta^* I - G_F\left(\zeta\right)' G_F\left(\zeta\right)\right] = D\left(\zeta\right)' D\left(\zeta\right) \tag{7.5.7}$$

where D is rational in ζ , has no poles inside or on the circle Γ , is invertible inside Γ , and the matrix coefficients of its power series expansion inside Γ can

 $^{^{8\,}}$ Appendix B displays a linear prediction problem that leads to the spectral factorization problem here.

be chosen to be real. The matrix analytic function D is unique only up to premultiplication by an orthogonal matrix but can be chosen to be independent of w_0 . The existence of this factorization follows from results about the linear extrapolation of covariance stationary stochastic processes. In particular, it is known from Theorems 4.2, 6.2 and 6.3 of Rozanov (1967) that the infimum of the objective is given by:

$$w_0' D(0)' D(0) w_0. (7.5.8)$$

When $\theta^* I - G'_F G_F$ is nonsingular on Γ , the infimum is attained. To verify this, write the first-order conditions for maximizing (7.5.4) subject to (7.5.5) and (7.5.6) as

$$\left[\theta^* I - G_F\left(\zeta\right)' G_F\left(\zeta\right)\right] W\left(\zeta\right) = \mathcal{L}\left(\zeta\right)',\tag{7.5.9}$$

where \mathcal{L} is the Lagrange multiplier on (7.5.5) and (7.5.6). Then the matrix D in the factorization (7.5.7) is nonsingular with an inverse that is rational and well defined on and inside the circle Γ . Substituting the factorization (7.5.7) into (7.5.9) gives

$$D(\zeta)' D(\zeta) W(\zeta) = \mathcal{L}(\zeta)', \qquad (7.5.10)$$

where $D(\zeta), W(\zeta)$, being analytic inside Γ , have expansions in nonnegative powers of ζ , and $D(\zeta)'$ and $\mathcal{L}(\zeta)'$ have expansions in nonpositive powers of ζ in the interior of Γ . If $D(\zeta)'$ is invertible, then following Whittle (1983, p. 100), $W(\zeta)$ satisfies

$$D(\zeta) W(\zeta) = \left[D(\zeta)'^{-1} \mathcal{L}(\zeta)' \right]_{+}$$

where $[\cdot]_+$ is the annihilation operator that sets negative powers of ζ to zero. Because $D(\zeta)'^{-1}$ and $\mathcal{L}(\zeta)'$ are both one-sided in nonpositive powers of ζ , $[D(\zeta)'^{-1}\mathcal{L}(\zeta)']_+ = D(0)'^{-1}\mathcal{L}(0)'$. Therefore, the solution is

$$D(\zeta) W(\zeta) = D(0)^{\prime - 1} \mathcal{L}(0)^{\prime}.$$
(7.5.11)

Then from (7.5.10), $\mathcal{L}(0)' = D(0)'D(0)W(0)$. Substituting into (7.5.11) gives

$$D(\zeta) W(\zeta) = D(0) w_0.$$
 (7.5.12)

In addition, the infimum is attained by 9

$$W^{*}(\zeta) = D(\zeta)^{-1} D(0) w_{0}. \qquad (7.5.13)$$

⁹ The factorization is also the key for calculating the projection of y_t on the semi-infinite history $x_s, s \leq t$ where $\{y_t, x_t\}$ is a covariance stationary process (see Whittle (1983, pp. 99-100)). Condition (7.5.10) corresponds to the solution of Whittle's projection problem where $D(\zeta)'D(\zeta)$ is interpreted as the spectral density of x and $\mathcal{L}(\zeta)$ is interpreted as the cross-spectral density between y and x.

190

Substituting into (7.5.4) confirms that the minimized solution is (7.5.8).

As is evident from the proof, the infimum in (7.5.4) may not be attained when $\theta^*I - G'_F G_F$ is singular somewhere on Γ . But this problem can be remedied by enlarging the space from \mathcal{W} to \mathcal{W}^a .

Corollary 7.5.1. Assume that F is such that $\int_{\Gamma} \log \det(\theta I - G'_F G_F) d\lambda > -\infty$. Then the problem

$$\min_{\mathcal{W}^{a}} \int_{\Gamma} W\left(\zeta\right)' \left[\theta^{*}I - G_{F}\left(\zeta\right)' G_{F}\left(\zeta\right)\right] W\left(\zeta\right) d\lambda\left(\zeta\right)$$

has a solution and the minimized value is $w'_0 D(0)' D(0) w_0$.

Proof. Solution (7.5.13) is in \mathcal{W}^a even when $\theta^* I - G'_F G_F$ is singular somewhere on Γ .

Corollary 7.5.1 shows that a solution exists for the multiplier problem, provided that the entropy restriction (7.5.2) is satisfied. But unless the matrix $(\theta^*I - G'_F G_F)$ is nonsingular at all frequencies, the minimizing shock sequence may not be stable and may not stabilize the state vector sequence. Problems occur when $W^*(\zeta) = D(\zeta)^{-1}D(0)w_0$ has a pole on Γ , or equivalently when $D(\zeta)^{-1}$ has a pole on Γ that is not annihilated by $D(0)w_0$. Nevertheless, even these destabilizing solutions for W^* can be approximated by a sequence of W's, each of which is in \mathcal{W} and hence each of which stabilizes the state vector sequence.

The multiplier problem depends on the initial condition $W(0) = w_0$. We now seek to replace this multiplier problem by an entropy criterion that does not depend on the initial condition. To accomplish this, we will eventually have to show that for a given θ^* , the control law that solves the multiplier game does not depend on the choice of initial condition w_0 and is the same control law that solves the entropy control problem. We shall do this in section 7.5.4.

The entropy criterion is motivated by the following representation:

Theorem 7.5.3. Assume that θ^* and F are such that $\int_{\Gamma} \log \det(\theta^* I - G'_F G_F) d\lambda > -\infty$. The criterion $\log \det[D(0)'D(0)]$ can be represented

$$\log \det \left[D\left(0\right)' D\left(0\right) \right] = \int_{\Gamma} \log \det \left[\theta^* I - G_F\left(\zeta\right)' G_F\left(\zeta\right) \right] d\lambda\left(\zeta\right).$$
(7.5.14)

Proof. D(0)'D(0) can be regarded as a 'one-step' prediction error covariance matrix for a vector process $D(L)\epsilon_t$, where L is the lag operator and ϵ_t is an

i.i.d. random process with mean zero and identity contemporaneous covariance matrix, and $D(\zeta)$ originates in the spectral factorization (7.5.7). We can use a result from linear prediction theory to verify the representation (7.5.14). See Theorem 6.2 of Rozanov (1967, page 76).

Theorem 7.5.2 and Theorem 7.5.3 both require that $\int_{\Gamma} \log \det(\theta^* I - G'_F G_F) d\lambda > -\infty$ but permit $\theta^* I - G'_F G_F$ to be singular at isolated points in Γ .

Evaluating the right-hand side of (7.5.14) requires no spectral factorization, just integration over frequencies. The contour integral on the right side of (7.5.14) is the entropy criterion. In the undiscounted case, it coincides with the measure of entropy used by Mustafa and Glover (1988).¹⁰ When $\beta = 1$, the *F* that maximizes (7.5.14) is often motivated as an approximation of the *F* that maximizes the H_{∞} criterion, one that maintains analyticity of *W*.

Next we show that when W^* stabilizes the state vector sequence, w_{t+1} can be represented as a function of the time t state x_t .

Theorem 7.5.4. Assume θ^* and F are such that $\theta^*I - G'_F G_F$ is nonsingular on Γ . Then the solution to the multiplier problem can be represented recursively as

$$w_{t+1} = K x_t \tag{7.5.15}$$

where

$$K = (\theta^* I - C' P C)^{-1} C' P A_F, \qquad (7.5.16)$$

and P is the positive semidefinite solution to the Riccati equation

$$P = H'_F H_F + \beta A'_F P A_F + \beta A'_F P C \left(\theta^* I - C' P C\right)^{-1} C' P A_F$$
(7.5.17)

for which $A_F + CK$ has eigenvalues that are inside the circle Γ . Moreover,

$$\int_{\Gamma} \log \det \left[\theta^* I - G'_F G_F\right] d\lambda = \log \det \left(\theta^* I - C' P C\right).$$
(7.5.18)

Proof. We use a recursive formulation and solution of the spectral factorization problem (7.5.7) to prove the theorem. To compute D in the spectral factorization $I\theta^* - G'_F G_F = D'D$, we apply the factorization result given by Zhou,

¹⁰ It coincides with their measure of entropy at $s_0 = \infty$.

Doyle, and Glover (1996). Recall that $G_F = H_F (I - \zeta A_F)^{-1} C$. The spectral density matrix to be factored is:

$$\theta^* I - G'_F G_F = \theta^* I - C' \left[I - \sqrt{\beta} \exp\left(-i\omega\right) A'_F \right]^{-1} H'_F H_F \left[I - \sqrt{\beta} \exp\left(i\omega\right) A_F \right]^{-1} C$$
$$= \theta^* I - C' \left[\exp\left(i\omega\right) I - \sqrt{\beta} A'_F \right]^{-1} H'_F H_F \left[\exp\left(-i\omega\right) I - \sqrt{\beta} A_F \right]^{-1} C,$$

where we have used the parameterization: $\zeta = \sqrt{\beta} \exp(i\omega)$. From Theorem 21.26 of Zhou, Doyle, and Glover (1996, pages 557 and 558), we obtain the factorization:

$$\theta^* I - C' \left[\exp(i\omega) I - \sqrt{\beta} A'_F \right]^{-1} H'_F H_F \left[\exp(-i\omega) I - \sqrt{\beta} A_F \right]^{-1} C$$

= $\left(I - C' \left[\exp(i\omega) I - \sqrt{\beta} A'_F \right]^{-1} \sqrt{\beta} K' \right) R \left(I - \sqrt{\beta} K \left[\exp(-i\omega) I - \sqrt{\beta} A_F \right]^{-1} C \right)$
= $\left(I - \zeta' C' \left[I - \zeta' A'_F \right]^{-1} K' \right) R \left(I - \zeta K \left[I - \zeta A_F \right]^{-1} C \right)$
(7.5.19)

where

192

$$R = \theta^* I - C' P C, \tag{7.5.20}$$

$$K = R^{-1} C' P A_F, (7.5.21)$$

and $P \geq 0$ is the stabilizing solution of the Riccati equation

$$\beta A'_F P \left(I - \frac{1}{\theta^*} C C' P \right)^{-1} A_F - P + H'_F H_F = 0.$$
 (7.5.22)

We establish that formula (7.5.22) is equivalent with (7.5.17) by showing that

$$\left(I - \frac{1}{\theta^*}CC'P\right)^{-1} = I + C\left(\theta^*I - C'PC\right)^{-1}C'P.$$

We verify this result by post multiplying the matrix $I - \frac{1}{\theta^*}CC'P$ by the matrix $I + C(\theta^*I - C'PC)^{-1}C'P$:

$$\left(I - \frac{1}{\theta^*}CC'P\right) \left[I + C\left(\theta I - C'PC\right)^{-1}C'P\right]$$

= $I - \frac{1}{\theta^*}CC'P + C\left(I - \frac{1}{\theta^*}C'PC\right)\left(\theta I - C'PC\right)^{-1}C'P$
= $I - \frac{1}{\theta^*}CC'P + \frac{1}{\theta^*}CC'P$
= I .

For the stabilizing solution, K from (7.5.21) is such that $I - \zeta \sqrt{\beta} K [I - \sqrt{\beta} \zeta A_F]^{-1} C$ has zeros outside the unit circle of the complex plane (Zhou, Doyle, and Glover (1996)). As a consequence, $I - \zeta K [I - \zeta A_F] C$ has zeros outside of the circle Γ . Therefore, (7.5.19) and (7.5.7) imply that

$$D^{*}(\zeta) = R^{1/2} \left(I - \zeta K \left[I - \zeta A_{F} \right]^{-1} C \right)$$
(7.5.23)

has zeros outside Γ , and

$$\theta^*I - G'_F G_F = D^{*'}D^*.$$

Furthermore,

$$D^{*}(0)' D^{*}(0) = R = \theta^{*}I - C'PC.$$

The entropy criterion (7.4.4) can thus be represented as $\log \det(\theta^* I - C'PC)$. From formula (7.5.12), the solution for $W(\zeta)$ can be represented as

$$D^{*}(\zeta) W(\zeta) = D^{*}(0) w_{0}.$$

Using (7.5.23) gives

$$\zeta^{-1} [W(\zeta) - w_0] = K (I - \zeta A_F)^{-1} CW(\zeta)$$

and using $X(\zeta) = (I - \zeta A_F)^{-1} CW(\zeta)$ gives the recursive formula

$$w_{t+1} = Kx_t.$$

Theorem 7.5.4 can be extended to allow for isolated singularities. In Appendix E we show that the entropy formula (7.5.18) of Theorem 7.5.4 continues to hold if $\theta^*I - G'_F G_F$ is positive semidefinite and nonsingular at either $\sqrt{\beta}$ or $-\sqrt{\beta}$.

Formula (@**Eq.riccati2**@) can also be written $P = H'_F H_F + A'_F \mathcal{D}(P) A_F = S(P)$ where the operators \mathcal{D} and S are defined in (6.2.7c) and (6.2.7f) on page 141. Note also that (7.5.16) matches (6.2.7e).

7.6. Relation of Stackelberg multiplier game to entropy criterion

One step remains to show that the Stackelberg multiplier game justifies the entropy criterion (7.4.4). The extra step is needed because criterion (7.5.8) depends on w_0 while (7.4.4) does not. But Theorem 6.6.1 showed that the F that solves (7.4.3) is independent of w_0 . Therefore, we will attain the same decision rule F by maximizing a criterion defined in terms of D(0)'D(0) alone, ignoring w_0 . Thus, let $w'_0\hat{D}(0)'\hat{D}(0)w_0$ denote criterion (7.5.8) for another control law, say \hat{F} . If

$$w_0' D(0)' D(0) w_0 \ge w_0' \hat{D}(0)' \hat{D}(0) w_0$$

for all w_0 , then

$$D(0)' D(0) \ge \hat{D}(0)' \hat{D}(0)$$

where \geq is the standard partial ordering of positive semidefinite matrices. As a consequence,

trace
$$\left[D\left(0\right)' D\left(0\right)\right] \ge \operatorname{trace}\left[\hat{D}\left(0\right)' \hat{D}\left(0\right)\right],$$

or alternatively

$$\log \det \left[D\left(0\right)' D\left(0\right) \right] \ge \log \det \left[\hat{D}\left(0\right)' \hat{D}\left(0\right) \right].$$

Because we want the criterion to apply for all initial conditions w_0 , we take our criterion to be

$$\log \det D\left(0\right)' D\left(0\right).$$

Theorem 7.5.3 shows that this is the entropy criterion used to define (7.4.4).

7.7. Etymology of 'entropy'

The criterion (7.4.4) acquires the name 'entropy' via formula (7.5.14), which links (7.4.4) to the log det of a one-step ahead prediction error covariance matrix for a process with moving average representation $D(L)\varepsilon_t$, where ε_t is an i.i.d. process with mean zero and identity covariance matrix. For a filtering problem, we also applied the term entropy to a closely related criterion that appears in (4.6.9) and (4.6.15) on pages 113 and 114, respectively. There the connection to a prediction problem was immediate, but here it is only indirect via the link revealed in formula (7.5.14) and the arguments in the proof of Theorem 7.5.2.¹¹

7.8. Risk sensitivity

This section explores the connection between the entropy criterion (7.5.18) and the discounted risk-sensitive criterion described by Hansen and Sargent (1995). Hansen and Sargent consider a situation where a decision maker is interested in evaluating fixed rules $u_t = -Fx_t$ from the point of view of minimizing a cost-criterion defined recursively as

$$C(x) = x'H'_FH_Fx + \beta \mathcal{R}(C(y)|x)$$
(7.8.1)

where

$$\mathcal{R}\left(\Upsilon|x\right) = -\left(\frac{2}{\sigma}\right)\log E\left(\exp\left(\frac{-\sigma\Upsilon}{2}\right)\Big|x\right),\tag{7.8.2}$$

where $H_F = H - JF$ and

$$y = A_F x + C w \tag{7.8.3}$$

where w is now an i.i.d. Gaussian sequence with mean zero and covariance matrix I. In (7.8.2), σ is the 'risk-sensitivity' parameter. When $\sigma < 0$, \mathcal{R} adds an additional aversion to risk beyond that embodied in the cost function C(y).

Define the operator 12

$$\mathcal{D}(V) = V - \sigma V C \left(I + \sigma C' V C\right)^{-1} C' V.$$
(7.8.4)

¹¹ Also, the presence of discounting compels us to use the change of measure associated with λ to reveal the connection to the log det of what looks like a prediction error covariance matrix.

 $^{^{12}\,}$ This operator also appears in chapters 2 and 6.

Drawing on results of Jacobson (1973), Hansen and Sargent (1995) show that

$$\mathcal{R}\left(\Upsilon|x\right) = x'A'_F\mathcal{D}\left(\Upsilon\right)A_Fx\tag{7.8.5}$$

and that the cost function C(x) that is the fixed point of (7.8.1) has the form

$$C(x) = x'V^*x + c^* (7.8.6)$$

where V^* is the fixed point of recursions on the operator

$$\mathcal{S}(V) = H'_F H_F + \beta A'_F \mathcal{D}(V) A_F \tag{7.8.7}$$

and

$$c^* = \frac{\beta}{1-\beta} \frac{1}{\sigma} \log \det \left(I + \sigma C' V^* C\right).$$
(7.8.8)

The connection between risk-sensitive preferences and a preference for robustness can be seen by letting $\sigma = -\theta^{-1}$ and noting that

$$\mathcal{D}(V) = V + \theta^{-1} V C \left(I - \theta^{-1} C' V C \right)^{-1} C' V$$
(7.8.9)

which makes (7.8.7) the same operator that appears on the right of (7.5.17). Also, c^* is

$$c^{*} = -\frac{1}{\theta} \frac{\beta}{1-\beta} \log \det \left(I - \theta^{-1} C' V^{*} C \right).$$
 (7.8.10)

This can also be written

$$c^* = -\frac{1}{\theta} \frac{\beta}{1-\beta} \left[-n\log\theta + \log\det\left(\theta I - C'V^*C\right) \right].$$
(7.8.11)

Consider representation (7.8.6) for the cost function. Here $x'V^*x$ is the part of the cost function contributed by the initial condition, while the log det term is contributed by the stochastic steady state. Consider minimizing the cost function starting from x = 0, so that only the log det term is relevant. The log det($\theta I - C'VC$) term can be interpreted as the log determinant of a onestep-ahead prediction error covariance matrix and so can be expressed as the right side of (7.5.14) for *some* stationary process with a particular associated spectral density matrix. From (7.5.17) and the definitions of the S(V) and $\mathcal{D}(V)$ operators, it follows that the appropriate spectral density is identical with that used in defining (7.5.14). In the case that x = 0, minimizing cost amounts to minimizing c^* . Because $\theta > 0$, this comes down to maximizing

$$\log \det \left(\theta I - C' V^* C\right). \tag{7.8.12}$$

This is equivalent with maximizing entropy defined by (7.5.14).

7.9. Risk aversion across frequencies

This section shows how the entropy criterion adjusts the H_2 criterion to express a concern about model misspecification by putting additional concavity into a utility function. We thereby develop a sense in which the entropy criterion represents model misspecification by inducing *risk aversion* across frequencies.

The H_2 criterion is

$$H_{2} = -\int_{\Gamma} \operatorname{trace} \left[G_{F}\left(\zeta\right)' G_{F}\left(\zeta\right) \right] d\lambda\left(\zeta\right),$$

and the entropy criterion is

$$\operatorname{ent} = \int_{\Gamma} \log \det \left[\theta I - G_F(\zeta)' G_F(\zeta) \right] d\lambda(\zeta) \,.$$

Take a symmetric negative semidefinite matrix V with eigenvalues $-\delta_1, \ldots, -\delta_n$. Let $\theta > \max_i -\delta_i$. Then trace $(V) = \sum_j -\delta_j$ and

$$\log \det (\theta I + V) = \sum_{j} \log (\theta - \delta_j).$$

Note that $\log(\theta - \delta)$ is a concave function of $-\delta$.

Associated with each ζ is a set of eigenvalues of $G_F(\zeta)'G_F(\zeta)$ that we denote $\delta_1(\zeta), \ldots, \delta_n(\zeta)$. Let them be ordered according to their magnitude. Then we can write the H_2 criterion as

$$H_{2} = \sum_{j} \int_{\Gamma} -\delta_{j} \left(\zeta\right) d\lambda \left(\zeta\right).$$

The entropy criterion is formed from H_2 be putting a concave transformation inside the integration:

$$\operatorname{ent} = \sum_{j} \int_{\Gamma} \log \left[\theta - \delta_{j} \left(\zeta \right) \right] d\lambda \left(\zeta \right).$$
(7.9.1)

Thus the entropy criterion puts more curvature into the return function. This has effects that could also be represented as enhanced risk aversion. Notice that here the 'risk aversion' seems to be across frequencies: in (7.9.1) we average over eigenvalues and frequencies instead of states of nature. Big eigenvalues have relatively more weight in the entropy criterion because of the concavity of the log function.

7.10. Concluding remarks

The decision maker's approximating model asserts that the Fourier transform of a target vector $Z(\zeta)$ is

$$Z\left(\zeta\right) = G_F\left(\zeta\right)w_0$$

where $G_F(\zeta)$ is the transfer function $G_F(\zeta) = H_F(I - (A - BF)\zeta)^{-1}C$ and F is the decision maker's feedback rule. The approximating model sets $W(\zeta) = w_0$, but the misspecified models assert that

$$Z\left(\zeta\right) = G_F\left(\zeta\right) W\left(\zeta\right).$$

Deviations of $W(\zeta)$ from w_0 represent the approximating model's misspecification of the temporal properties of the shock process.¹³

Without fear of model misspecification, the decision maker would choose F to maximize H_2 defined in equation (7.3.6). A preference for robustness to model misspecification can be expressed by having the decision maker replace H_2 by either H_∞ or an entropy criterion. The H_∞ criterion induces a robust decision rule via the following thought process. The decision maker considers perturbations to the temporal properties of the shocks and wants decisions that will work well across a broad set of such patterns. To promote robustness, the decision maker investigates the consequences of his rule under the worst shock process. But what is worst depends on his decision rule. Given his decision rule, the worst serial correlation pattern focuses spectral power at the frequency that attains the highest weight in the frequency domain representation of $Z(\zeta)'Z(\zeta)$. The contribution of that frequency to discounted costs is measured by the maximal eigenvalue of $G_F(\zeta)'G_F(\zeta)$. The decision maker achieves a robust rule by optimizing against that worst serial correlation pattern, in particular by selecting the feedback rule that minimizes the maximum eigenvalue across all frequencies. Under the entropy criterion the decision maker responds in a similar but less severe way by flattening the response $G_F(\zeta)$ across ζ 's. We study an example of such behavior in chapter 9, where we use such insights from the frequency domain to interpret how a form of precautionary savings is called for by a robust decision rule for a permanent income model.

¹³ See appendix E for an interpretation of $W(\zeta)$ in terms of the spectral density matrix of a random vector of shocks.

A. Infimization of H_{∞}

To verify that we have found the infimum of version 2 of (7.3.7)-(7.3.8), let ω^* be the frequency associated with the maximum value of ρ and let $v(\omega^*)$ denote the corresponding eigenvector. This eigenvector can be complex. We can find a $W^*(\zeta)$ with all real coefficients, with an initial coefficient zero, and that coincides with $v(\omega^*)$ for $\zeta = \sqrt{\beta} \exp(i\omega^*)$. We accomplish this while setting all values of w_t to zero except possibly those for w_1 and w_2 . In particular, that the coefficients of $W^*(\zeta)$ be real requires symmetry, i.e., $W^*(\sqrt{\beta} \exp(i\omega))' = W^*(\sqrt{\beta} \exp(-i\omega))^\top$, where $\overleftarrow{}$ denotes transposition. This leads to two equations of the form $W^*(\zeta^*) = w_1\zeta^* + w_2\zeta^{*2}$, $W^*(\zeta^{*'}) = w_1 \zeta^{*'} + w_2 \zeta^{*'^2}$, where here ' denotes the complex conjugate, and $\zeta^* =$ $\sqrt{\beta} \exp(i\omega)$. These two equations determine real valued vectors w_1, w_2 . To form the infinizing $W(\zeta)$, we shall construct an approximating sequence of 'distributed lags' of $W^*(\zeta)$ that converge to it. To get distributed lags of the desired form, create a sequence of continuous positive scalar functions $\{g_n\}$ such that:

(i)
$$g_n(\omega) = g_n(-\omega);$$

(ii) $\frac{1}{2\pi} \int_{-\pi}^{\pi} g_n(\omega) d\omega = 1$

(ii)
$$\frac{1}{2\pi} \int_{-\pi}^{\pi} g_n(\omega) d\omega = 1;$$

- (iii) $\{g_n(\omega^*)\}$ diverges; (iv) $\{g_n\}$ converges uniformly to zero outside any open interval containing ω^* ;
- (v) $\int_{-\pi}^{\pi} \log g_n(\omega) d\omega > 0.$

Then associated with each g_n is a real scalar (one-sided) sequence with transform $b_n(\zeta)$ such that $b_n(\zeta)^* b_n(\zeta) = g_n(\omega)$ for $\zeta = \sqrt{\beta} \exp(i\omega)$.

Construct $W_n(\zeta) \propto b_n(\zeta) W^*(\zeta)$, where the constant of proportionality makes the resulting W_n satisfy constraint (7.3.8). We have designed the sequence $\{W_n\}$ to approximate the direction $v(\omega^*)$. The sequence of transforms $\{g_n\}$ converges to a generalized function, namely a Dirac-delta function with mass concentrated at frequency ω^* . It is straightforward to show that:

$$\lim_{n \to \infty} \int_{\Gamma} W_n(\zeta)' G_F(\zeta)' G_F(\zeta) W_n(\zeta) d\lambda(\zeta) = \eta (H_{\infty})^2.$$

B. A dual prediction problem

A prediction problem is dual to maximizing (7.5.4) subject to (7.5.5)–(7.5.6). Let $[\theta I - G_F(\zeta)'G_F(\zeta)]$ for $\zeta = \sqrt{\beta} \exp(i\omega)$ denote a spectral density matrix for a covariance stationary process $\{y_t\}$. The purpose is to predict $(w_0)'y_t$ linearly from past values of y_t . A candidate forecast rule of the form

$$-\sum_{j=1}^{\infty} \left(w_j\right)' y_{t-j} \tag{7.B.1}$$

has forecast error

$$\sum_{j=0}^{\infty} \left(w_j \right)' y_{t-j}$$

Then criterion (7.5.4) is interpretable as the forecast-error variance associated with this prediction problem. The constraints (7.5.6) prevent the forecast from depending on y_{t+j} for $j \ge 1$.

C. Duality

7.C.1. Evaluating a given control law

For a given control law F form the corresponding G_F and define:

$$\theta_F = H^2_{\infty}\left(F\right).$$

It follows that for all $W(\zeta)$

$$\theta_F \int_{\Gamma} W' W d\lambda \geq \int_{\Gamma} W' G'_F G_F W d\lambda.$$

Therefore, for all $\theta \geq \theta_F$, $\int_{\Gamma} W' \left[\theta I - G'_F G_F \right] W d\lambda$ is well defined for all $\theta \geq \theta_F$ but not for $\theta < \theta_F$.

For fixed F, consider the inf part of Game 2 (7.3.7):

Original (Worst Case) Minimization Problem Lars: let's double check that we mean W and not W^a as the set over which we min. See the definitions on page 182.

Problem 1
$$\min_{\mathcal{W}} - \int_{\Gamma} W' G'_F G_F W d\lambda$$

200

Duality

subject to

$$\int_{\Gamma} W'Wd\lambda \leq w_0'w_0 + \eta$$

This problem minimizes a concave function subject to a convex constraint set, so standard duality theory does not apply. In the interests of applying duality theory, we study the following alternative problem:

A Related Constrained Problem:

Problem 2
$$\min_{\mathcal{W}} \int_{\Gamma} W' \left(\theta_F I - G'_F G_F \right) W d\lambda$$

subject to:

$$\int_{\Gamma} W' W d\lambda \le \eta + {w_0}' w_0.$$

This problem is to minimize a convex function subject to a convex constraint set, so duality theory applies to it. We shall first show that a solution of Problem 2 with binding constraint also solves Problem 1. Then we shall apply standard duality theory to problem 2.

Theorem 7.C.1. A solution to problem 2 with binding constraint solves problem 1. *Proof.* Let W^* solve Problem 2 with the magnitude constraint binding:

$$\int_{\Gamma} W^{*\prime} W^{*} d\lambda = \eta + w_0^{\prime} w_0$$

and

$$W^{*}\left(0\right)=w_{0}.$$

Consider any other W such that

$$\int_{\Gamma} W' W d\lambda \le \eta + w_0' w_0.$$

 $\quad \text{and} \quad$

$$W\left(0\right)=w_{0}.$$

Then

$$\int_{\Gamma} W'\left(\theta_{F}I - G'_{F}G_{F}\right) W d\lambda \geq \int_{\Gamma} W^{*'}\left(\theta_{F}I - G'_{F}G_{F}\right) W^{*} d\lambda,$$

and

$$\theta_F \int_{\Gamma} W' W d\lambda \leq \theta_F \int_{\Gamma} W^{*'} W^* d\lambda.$$

Therefore

$$-\int_{\Gamma} W'G'_F G_F W d\lambda \ge -\int_{\Gamma} W^{*'}G'_F G_F W^* d\lambda,$$

which implies that W^* also solves Problem 1.

Thus a way to solve Problem 1 is to solve Problem 2 and verify the solution satisfies the magnitude constraint with equality.

We now apply duality theory to problem 2 by forming:

Saddle Point Version of Problem 2:

$$\inf_{\mathcal{W}} \sup_{\theta \ge \theta_F} \left[\int_{\Gamma} W' \left(\theta I - G'_F G_F \right) W d\lambda - \left(\theta - \theta_F \right) \left(\eta + w'_0 w_0 \right) \right].$$

We interpret $\theta - \theta_F$ as the Lagrange multiplier for Problem 2 and θ as the Lagrange multiplier for Problem 1. Because Problem 2 entails minimizing a convex function subject to a convex constraint set, standard duality theory applies to it. The conjugate problem is obtained by switching the order of the inf and sup operations:

$$\sup_{\theta \ge \theta_F} \inf_{\mathcal{W}} \left[\int_{\Gamma} W' \left(\theta I - G'_F G_F \right) W d\lambda - \left(\theta - \theta_F \right) \left(\eta + w_0' w_0 \right) \right].$$
(7.C.1)

We can use this problem to construct the Lagrange multiplier θ for each $\eta > 0$.

By construction the saddle-point value for the conjugate problem coincides with the optimized value for Problem 2. When the specification-error constraint is binding for Problem 2, we can obtain the optimized value for Problem 1 by subtracting the constant $\theta_F(\eta + w'_0 w_0)$ from (7.C.1). The resulting conjugate problem is

$$\sup_{\theta \ge \theta_F} \inf_{\mathcal{W}} \left[\int_{\Gamma} W' \left(\theta I - G'_F G_F \right) W d\lambda - \theta \left(\eta + w'_0 w_0 \right) \right].$$
(7.C.2)

Thus we have eliminated the influence of θ_F on the objective of the saddle-point problem. But θ_F still affects the constraint set limiting the choice of θ (through the appearance of θ_F under the sup operator). This dependence can also be removed by virtue of the following theorem.

Theorem 7.C.2. If the value of (7.C.2) is finite, then $\theta \ge \theta_F$.

Proof. Suppose that $\theta < \theta_F$, and consider the inner infimum part of the saddle-point problem (7.C.2):

$$\inf_{\mathcal{W}} \int_{\Gamma} W' \left(\theta I - G'_F G_F \right) W d\lambda.$$
(7.C.3)

Given the construction of θ_F , $(\theta I - G'_F G_F)$ has negative eigenvalues for some $|\zeta^*| = \sqrt{\beta}$. Parameterize Γ by forming $\zeta = \sqrt{\beta} \exp(i\omega)$, and let ω^* be the frequency associated with ζ^* . Thus there exists a complex vector v such that

$$v'\left(\theta I - G'_F G_F\right)v < 0$$

202

on a nondegenerate interval of ω 's containing ω^* . Imitating the argument in Appendix A, we can form a $W^*(\zeta) = w_1\zeta + w_2\zeta$ such that $W^*(\zeta^*) = v$. We can then use the Appendix A construction to form: $W_n(\zeta) \sim b_n(\zeta) W^*(\zeta)$. Then it is straightforward to show that:

$$\lim_{n \to \infty} \int_{\Gamma} W'_n \left(\theta I - G'_F G_F \right) W_n d\lambda = v' \left[\theta I - G_F \left(\zeta^* \right)' G_F \left(\zeta^* \right) \right] v < 0.$$

By construction $W_n(0) = 0$ and hence fails to satisfy the constraint for problem (7.C.3). Also problem (7.C.3) does not constrain the magnitude of W. We now form the sequence:

$$W_n = nW_n + w_0,$$

which by construction satisfies $\tilde{W}_n(0) = w_0$. Given our multiplication of W_n by n, it clearly follows that

$$\lim_{n \to \infty} \int_{\Gamma} W'_n \left(\theta I - G'_F G_F \right) W_n d\lambda = -\infty.$$

Therefore, the optimized value of problem (7.C.3) is $-\infty$ whenever $\theta < \theta_F$.

Given what the theorem establishes about the behavior of the inner infimum part of saddle-point problem (7.C.2) when $\theta < \theta_F$, we can state that (7.C.2) equals (7.C.3) defined as:

Conjugate Saddle Point Version of Problem 1

$$\sup_{\theta} \inf_{\mathcal{W}} \left[\int_{\Gamma} W' \left(\theta I - G'_F G_F \right) W d\lambda - \theta \left(\eta + w'_0 w_0 \right) \right].$$
(7.C.4)

Whenever this problem has a solution for W that satisfies the specification-error constraint with equality, the resulting W also solves Problem 1 and the value of the conjugate saddle-point problem coincides with that of Problem 1. This conjugate problem provides the Lagrange multiplier $\theta \ge \theta_F$ associated with Problem 1. Armed with this multiplier, consider the inner infimum problem, which we call the *multiplier* problem:

(Problem 3)
$$\inf_{\mathcal{W}} \left[\int_{\Gamma} W' \left(\theta I - G'_F G_F \right) W d\lambda \right].$$

The solution of Problem 3 coincides with that of the prediction problem described in Appendix B and analyzed in the text.

Given any η , we have just shown how to find the multiplier θ . We now suppose that the multiplier $\theta > \theta_F$ is given and want to deduce the corresponding value of η . Thus, suppose that we have a solution of the multiplier problem (Problem 3). It is sufficient for this problem to have a solution with $\theta > \theta_F$. (Later we shall discuss the case in which $\theta = \theta_F$.) We assume that:

$$\int \log \det \left(\theta_F I - G'_F G_F\right) d\lambda > -\infty. \tag{7.C.5}$$

Later we will describe what happens when this condition is violated.

Theorem 7.C.3. Suppose that $\theta > \theta_F$ and that $W(\zeta)$ solves the multiplier Problem 3. Then there exists $\eta > 0$ such that $W(\zeta)$ solves Problem 1.

Proof. From the dual prediction problem of Appendix B, we know that when $\theta > \theta_F$, the solution to the multiplier problem is:

$$W(\zeta) = D(\zeta)^{-1} D(0) w_0$$
(7.C.6)

where

$$D'D = \left(\theta I - G'_F G_F\right)$$

and D is continuous and nonsingular on the region $|\zeta| \leq \sqrt{\beta}$. Notice that D depends implicitly on θ The resulting objective function is: $w'_0 D(0)' D(0) w_0$. The η corresponding to this choice of θ satisfies:

$$\eta = \int w_0' D(0)' \left(\theta I - G_F' G_F \right)^{-1} D(0) w_0 d\lambda - w_0' w_0.$$
 (7.C.7)

7. C.2. When $\theta = \theta_F$

Next consider the possibilities when θ is equal to the lower threshold value θ_F . Condition (7.C.5) implies that we can still obtain the factorization:

$$D'D = \theta_F I - G'_F G_F,$$

where D is nonsingular on the region $|\zeta| < \sqrt{\beta}$, but now it is singular at some points $|\zeta| = \sqrt{\beta}$. Thus the candidate solution for W given by (7.C.6) may not be well defined, and the infimum in the multiplier Problem 3 may not be attained. Nevertheless, the infimum is still given by the quadratic form: $w'_0 D(0)' D(0) w_0$ and the implied η_F satisfies (7.C.7), and will typically be infinite.

When $\eta_F = \infty$, we can find a $\theta > \theta_F$ that yields any positive η . Sometimes η_F is finite for a small (Lebesgue measure zero) set of initializations w_0 . When this happens, we may only find $\theta \ge \theta_F$ for values of $\eta \le \eta_F$.
7.C.3. Failure of entropy condition

Finally, we consider what happens when

$$\int \log \det \left(\theta_F I - G'_F G_F \right) d\lambda = -\infty$$

Since G_F is a rational function of ζ with no poles in the region $|\zeta| \leq \sqrt{\beta}$, $\theta_F I - G'_F G_F$ is singular for all $|\zeta| = \sqrt{\beta}$. Factorizations still exist now take the form:

$$D'D = \theta_F I - G'_F G_F$$

where D has fewer rows than columns and has full rank on the region $|\zeta| < \sqrt{\beta}$ (see Rozanov (1967) pages 43–50). This makes it possible to have a variety of solutions to Problem 2, including solutions for which the specification-error is slack.

To understand better the multiplicity, note that it is now possible to find a \bar{W} such that:

$$D\tilde{W} = 0 \tag{7.C.8}$$

and for which $\tilde{W}(0) = 0$. Given any solution W^* to Problem 2, we may form $W^* + r\tilde{W}$ for any real number r without altering the objective of Problem 2. The value of r is restrained by the specification-error constraint, but it possible for this range to be nondegenerate.

When the specification-error constraint for Problem 2 can be slack at an optimum, the Lagrange multiplier, $\theta - \theta_F$, is zero, or equivalently $\theta = \theta_F$. Problem 2 will then have solutions in which the specification-error constraint is binding (but with a zero multiplier), and it is *only* these solutions that also solve Problem 1. As a consequence, solving the multiplier problem (Problem 3) for choices of θ greater than θ_F may not correspond to fixing an η for Problem 1. We illustrate this possibility in the following example.

Exceptional Example

In this example, we construct a \tilde{W} satisfying (7.C.8) and $\tilde{W} > 0 \forall \zeta \in \Gamma$. Suppose that A - BF = 0 and hence $G_F = H_F C$, which is constant across frequencies. Then θ_F is the largest eigenvalue of the symmetric matrix $C'H'_FH_FC$, and $\det[\theta_F I - G'_FG_F] = 0$ for all $\zeta \in \Gamma$. Let μ be an eigenvector associated with θ_F with norm one. Solutions W^* to Problem 2 are given by:

$$w_0^* = w_0$$
$$w_t^* = \alpha_t \mu$$

for t > 0 and the real numbers α_t chosen so that the magnitude constraint is satisfied. The resulting objective for Problem 2 is:

$$w_0' \left(\theta_F I - C' H'_F H_F C \right) w_0.$$

Provided that $\eta > 0$, the magnitude constraint can be made slack (say by letting α_t be zero).

A solution to Problem 1 is obtained by setting α_t to make the magnitude constraint be satisfied with equality. Then the objective for Problem 1 is:

$$-\theta_F\eta - w_0'C'H'_FH_FCw_0$$

Finally, the Lagrange multiplier obtained from the conjugate problem is given by its lower threshold θ_F .

Optimizing the Control Law

We next study what happens when the control law is chosen among a family of admissible laws. The choice of F alters the transfer function G_F , and we are led to study the game:

$$\max_{\mathcal{F}} \inf_{\mathcal{W}} - \int_{\Gamma} W' G'_F G_F W d\lambda$$

subject to

206

$$\int_{\Gamma} W' W d\,\lambda \leq \eta + w'_0 w_0.$$

Again it is fruitful to analyze a conjugate formulation. With this in mind, first solve:

$$C(\theta, F) = \inf_{\mathcal{W}} \left[\int_{\Gamma} W' \left(\theta I - G'_F G_F \right) W \, d\lambda - \theta \left(\eta + w'_0 w_0 \right) \right]$$

for a given (θ, F) pair. Then solve the conjugate game:

$$\max_{\mathcal{F},\theta} C\left(\theta,F\right) = \max_{\mathcal{F}} \sup_{\theta} C\left(\theta,F\right) = \max_{\theta} \sup_{\mathcal{F}} C\left(\theta,F\right).$$

Therefore given a solution F^* to the original game we can find a corresponding θ^* such that (F^*, θ^*) solves the conjugate game. Moreover, if F^* is optimal for all nonzero initializations w_0 , then F^* solves the entropy criterion associated with this θ^* . We want to show the converse.

Theorem 7.C.4. Fix a θ^* . Find the F^* that solves the entropy problem for θ^* .

Compute $\hat{\theta} = H_{\infty}(F^*)^2$ and verify that the control law F^* satisfies:

$$\int_{\Gamma} \log \det \left(\hat{\theta} I - G_F^{*\prime} G_F^* \right) d\lambda > -\infty \tag{7.C.9}$$

where G_F^* is the transfer function associated with the control law F^* . Then there exists W^* and an $\eta^* > 0$ such that F^*, W^* solves Game 2.

Proof. If inequality (7.C.9) is satisfied, factor $\theta^* I - G_F^{*\prime} G_F^*$

$$\theta^* \left(I - G_F^{*'} G_F^* \right) = D^{*'} D^*,$$

and construct the W^* :

$$W^{*}(\zeta) = D^{*}(\zeta)^{-1} D^{*}(0) w_{0}.$$

Then find η^* that solves

$$\eta^* = \int_{\Gamma} W^{*\prime} W^* d\lambda - w_0^{\prime} w_0.$$

D. Proof of theorem 7.5.4

This appendix restates a version of Theorem 7.5.4 under weaker assumptions about the nonsingularity of $[\theta I - G_F(\zeta)'G_F(\zeta)]$.

Theorem 7.D.1. Suppose that

i A_F has eigenvalues that are inside the circle Γ ; ii $\theta I - G'_F G_F \ge 0$ on Γ ; iii Either $\theta I - G_F(-\sqrt{\beta})' G_F(-\sqrt{\beta})$ or $\theta I - G_F(\sqrt{\beta})' G_F(\sqrt{\beta})$ is nonsingular.

Then the $H_{\text{entropy}}(\theta)$ criterion can be represented as

$$\log \det D(0)' D(0) = \log \det (\theta I - C' PC)$$

where P is defined implicitly by equation (7.D.3) below.

Proof. We prove this theorem by referring to results from Zhou, Doyle, and Glover (1996). We outline the proof in four steps.

Step One: Transform the discrete discounted formulation into continuous undiscounted formulation. Suppose that $\theta I - G_F(-\sqrt{\beta})'G_F(-\sqrt{\beta})$ is nonsingular. Define the linear fractional transformation:

$$\zeta = -\sqrt{\beta} \left(\frac{s + \sqrt{\beta}}{s - \sqrt{\beta}} \right). \tag{7.D.1}$$

This transformation maps $s = -\sqrt{\beta}$ into $\zeta = 0$, s = 0 into $\sqrt{\beta}$, $s = \infty$ into $-\sqrt{\beta}$. The transformation maps the imaginary axis into the circle Γ and points on the left side of the complex plane into points inside the circle.

Note also that

$$\beta \zeta^{-1} = -\sqrt{\beta} \left(\frac{-s + \sqrt{\beta}}{-s - \sqrt{\beta}} \right).$$

In the case that $\theta I - G_F(\sqrt{\beta})'G_F(\sqrt{\beta})$ is singular, we replace linear fractional transformation (7.D.1) with:

$$\zeta = \sqrt{\beta} \left(\frac{s + \sqrt{\beta}}{s - \sqrt{\beta}} \right). \tag{7.D.2}$$

In what follows we will use (7.D.1) but the argument for (7.D.2) is entirely similar. Step Two: Use parameterization (7.D.1) to write:

$$G_F(\zeta) = \left(s - \sqrt{\beta}\right) H_F\left[\left(s - \sqrt{\beta}\right)I + \left(s + \sqrt{\beta}\right)\sqrt{\beta}A_F\right]^{-1}C$$

$$= \left(s - \sqrt{\beta}\right) H_F\left[s\left(I + \sqrt{\beta}A_F\right) - \sqrt{\beta}\left(I - \sqrt{\beta}A_F\right)\right]C$$

$$= \left(s - \sqrt{\beta}\right) H_F\left(sI - \hat{A}\right)^{-1}\hat{C}$$

$$= \hat{G}_F(s)$$

where

$$\hat{A} = \sqrt{\beta} \left(I + \sqrt{\beta} A_F \right)^{-1} \left(I - \sqrt{\beta} A_F \right)$$
$$\hat{C} = \left(I + \sqrt{\beta} A_F \right)^{-1} C$$

Rewrite \hat{G}_F as

$$\hat{G}_{F}(s) = sH_{F}(sI - \hat{A})^{-1}\hat{C} - \sqrt{\beta}H_{F}(sI - \hat{A})^{-1}\hat{C} = H_{F}(sI - \hat{A})(sI - \hat{A})^{-1}\hat{C} + H_{F}\hat{A}(sI - \hat{A})^{-1}\hat{C} - \sqrt{\beta}H_{F}(sI - \hat{A})^{-1}\hat{C} = H_{F}\hat{C} - \hat{H}_{F}(sI - \hat{A})^{-1}\hat{C},$$

where

$$\hat{H}_F = H_F\left(\sqrt{\beta}I - \hat{A}\right).$$

Notice that

$$H_F \hat{C} = H_F \left(I + \sqrt{\beta} A_F \right)^{-1} C = \hat{G}_F \left(\infty \right) = G_F \left(-\sqrt{\beta} \right).$$

Step Three: Write for s imaginary

$$\theta I - \hat{G}'_F \hat{G}_F = \left(\hat{C}' \left(-sI - \hat{A}' \right)^{-1} \quad I \right) \begin{pmatrix} -\hat{H}'_F \hat{H}_F & \hat{H}'_F H_F \hat{C} \\ \hat{C}' H'_F \hat{H}_F & \theta I - \hat{C}' H'_F H_F \hat{C} \end{pmatrix} \begin{pmatrix} \left(sI - \hat{A} \right)^{-1} \hat{C} \\ I \end{pmatrix}$$

Notice that

$$\theta I - \hat{C}' H'_F H_F \hat{C} = \theta I - G_F \left(-\sqrt{\beta} \right)' G_F \left(-\sqrt{\beta} \right)$$

is nonsingular and in fact positive definite.

Step Four: Apply Corollary 13.20 of Zhou, Glover, and Doyle (1996) to conclude that there exists a matrix F such that:

$$\theta I - \hat{G}'_F \hat{G}_F = \left[I - \hat{C}' \left(-sI - \hat{A}' \right)^{-1} F' \right] \left(\theta I - \hat{C}' H'_F H_F \hat{C} \right) \left[I - F \left(sI - \hat{A} \right)^{-1} \hat{C} \right].$$

Now inverse transform from s to $\zeta.$ The following are useful formulas for carrying out this transformation. First

$$\hat{A} = \sqrt{\beta} \left(I + \sqrt{\beta} A_F \right)^{-1} \left(I - \sqrt{\beta} A_F \right).$$

Invert this relation to find that:

$$\left(I + \sqrt{\beta}A_F\right)\hat{A} = \sqrt{\beta}I - \beta A_F$$
$$\sqrt{\beta}\left(\hat{A} + \sqrt{\beta}I\right)A_F = -\left(\hat{A} - \sqrt{\beta}I\right)$$

or

or

or

or

$$A_F = \frac{1}{\sqrt{\beta}} \left(\sqrt{\beta}I + \hat{A} \right)^{-1} \left(\sqrt{\beta}I - \hat{A} \right).$$

Similarly,

$$\left(s - \sqrt{\beta}\right)\zeta = -\sqrt{\beta}\left(s + \sqrt{\beta}\right)$$
$$\left(\zeta + \sqrt{\beta}\right)s = \sqrt{\beta}\left(\zeta - \sqrt{\beta}\right)$$
$$s = \sqrt{\beta}\left(\frac{\zeta - \sqrt{\beta}}{\zeta + \sqrt{\beta}}\right)$$

Write:

$$I - F(sI - \hat{A})^{-1} \hat{C} = I - (\zeta + \sqrt{\beta}) F\left[\sqrt{\beta} (\zeta - \sqrt{\beta}) I - (\zeta + \sqrt{\beta}) \hat{A}\right]^{-1} \hat{C}$$

$$= I - (\zeta + \sqrt{\beta}) F\left[\zeta (\sqrt{\beta}I - \hat{A}) - \sqrt{\beta} (\sqrt{\beta}I + \hat{A})\right]^{-1} \hat{C}$$

$$= I + (\zeta + \sqrt{\beta}) F(I - \zeta A_F)^{-1} \frac{1}{\sqrt{\beta}} (\sqrt{\beta}I + \hat{A})^{-1} \hat{C}$$

$$= I + \frac{(\zeta + \sqrt{\beta})}{2\sqrt{\beta}} F(I - \zeta A_F)^{-1} C$$

$$= \tilde{G}_F(\zeta).$$

Note that

$$I + \frac{\left(\zeta + \sqrt{\beta}\right)}{2\sqrt{\beta}}F\left(I - \zeta A_F\right)^{-1}C = I + \frac{1}{2}FC + \frac{\zeta}{2\sqrt{\beta}}F\left(I + \sqrt{\beta}A_F\right)\left(I - \zeta A\right)^{-1}C.$$

Define *P* implicitly by:

$$\theta I - C'PC = \left(I + \frac{1}{2}C'F'\right) \left[\theta I - C'\left(I + \sqrt{\beta}A'_F\right)^{-1}H'_FH_F\left(I + \sqrt{\beta}A_F\right)^{-1}C\right] \left(I + \frac{1}{2}FC\right).$$
(7.D.3)

E. Stochastic interpretation of H_2

This appendix displays another game that implies H_2 where the shocks w_t are permitted to be nonzero for t > 0. Recall that w_t is $m \times 1$, where m is the number of shocks. We continue to assume that $w_t = 0$ for all t < 0. We state

Game 1a: Choose $(F, \{w_t\})$ to attain

$$\max_{\mathcal{F}} \inf_{\{w_t\}} - \sum_{t=0}^{\infty} \beta^t z'_t z_t \tag{7.E.1}$$

subject to

$$x_0 = Cw_0 \tag{7.E.2a}$$

$$\sum_{t=0}^{\infty} \beta^t w_t w_t' = \sigma^2 I \tag{7.E.2b}$$

$$\sum_{t=0}^{\infty} \left(\beta^{\frac{t}{2}} w_t\right) \left(\beta^{\frac{t-j}{2}} w_{t-j}\right)' = 0 \ \forall j \neq 0 \tag{7.E.2c}$$

$$\sigma^2 \le \eta \tag{7.E.2d}$$

Equations (7.E.2b), (7.E.2c) imply that

$$W(\zeta)W(\zeta)' = \sigma^2 I, \quad |\zeta| = \sqrt{\beta}, \tag{7.E.3}$$

Further, (7.E.3) implies (7.E.2b), (7.E.2c).

Game 1a has the following counterpart in the frequency domain:

Game 1b: Find (F, σ^2) that attain

$$\max_{F} \inf_{\sigma^{2}} -\sigma^{2} \int_{\Gamma} \operatorname{trace} \left[G_{F}\left(\zeta\right)' G_{F}\left(\zeta\right) \right] d\lambda\left(\zeta\right), \qquad (7.E.4)$$

subject to

$$\sigma^2 \le \eta. \tag{7.E.5}$$

We have substituted (7.E.3) into (7.3.7) to obtain (7.E.4). The solution of game 1b sets σ^2 at its upper bound η , and sets F to maximize the H_2 criterion (7.3.6).

7.E.1. Stochastic counterpart

Criterion (7.3.6) emerges when the shock process $\{w_t\}_{t=1}^{\infty}$ is taken to be a martingale difference sequence adapted to J_t , the sigma algebra generated by x_0 and the history of w, where $Ew_{t+1}w'_{t+1}|J_t = I$. The martingale difference specification implies

$$E\sum_{t=0}^{\infty} \left(\beta^{\frac{t}{2}} w_t\right) \left(\beta^{\frac{t-j}{2}} w_{t-j}\right)' = \begin{cases} \sigma^2 \left(1-\beta\right)^{-1} I & \text{if } j=0;\\ 0 & \text{otherwise.} \end{cases}$$
(7.E.6)

Equation (7.E.6) is equivalent with $EW(\zeta)W(\zeta)' = \sigma^2(1-\beta)^{-1}I$ for $\zeta \in \Gamma$. With this representation, (7.3.6) is proportional to $-(1-\beta)^{-1}E\sum_{t=0}^{\infty}\beta^t z_t z'_t$.¹⁴

¹⁴ See Whiteman (1985b).

Chapter 8 Calibrating θ with detection probabilities

8.1. The role of randomness

Though we are really interested in random processes, most of our calculations have been cast in terms of deterministic models. This has been true even when we studied filtering problems. In omitting explicit mention of randomness, we have exploited the mathematical structure of models with quadratic objective functions, linear transition laws, and Gaussian disturbances. Control and filtering of such models involve, after mathematical expectations have been taken appropriately, deterministic manipulations of moment matrices. Thus, the certainty equivalence principle stated on page 27 implies that we would derive the same decision rules had we included i.i.d. Gaussian shocks in the models.

In this chapter we explicitly include randomness in order to characterize models that are difficult to distinguish from the approximating model using moderate amounts of data. The presence of randomness in the transition law conceals the distortion of the alternative model relative to the approximating model and makes it statistically difficult to detect if the distortion is not too big.

We use a statistical theory of detection to define a mapping from θ to a detection error probability for discriminating between the approximating model and the endogenous worst case model associated with that θ . We use that detection error probability to determine a context-specific θ that is associated with a set of alternative models against which it is reasonable to want to be robust.¹

¹ In the context of continuous time models, Anderson, Hansen, and Sargent (2001) investigate the connection among detection error probabilities, a preference for robustness, and alterations of market prices for risk.

8.1.1. Approximating and distorting models

For a given decision rule $u_t = -Fx_t$, we assume that the approximating model makes the state evolve according to the stochastic difference equation

$$x_{t+1} = A_o x_t + C\check{\epsilon}_{t+1}, \tag{8.1.1}$$

where now $\check{\epsilon}_{t+1}$ is an i.i.d. sequence of Gaussian disturbances with mean zero and identity contemporaneous covariance matrix. We'll represent a distorted model as

$$x_{t+1} = A_o x_t + C \left(\epsilon_{t+1} + w_{t+1} \right),$$

= $\hat{A} x_t + C \epsilon_{t+1}$ (8.1.2)

where $\hat{A} = A_o + C\kappa(\theta)$, $w_{t+1} = \kappa(\theta)x_t$, and ϵ_{t+1} is another i.i.d. Gaussian vector with mean 0 and identity covariance matrix. The transition densities associated with models (8.1.1) and (8.1.2) are absolutely continuous with respect to each other, i.e., they put positive probabilities on the same events.² Models that are not absolutely continuous with respect to each other are easy to distinguish empirically.

8.2. Detection error probabilities

Detection error probabilities can be calculated using likelihood ratio tests. Thus, consider two alternative models. Model A is the approximating model (8.1.1), and model B is the distorted model (8.1.2) associated with the context specific worst case shock implied by θ . Consider a fixed sample of observations on the state $x_t, t = 0, \ldots, T-1$. Let L_{ij} be the likelihood of that sample for model j assuming that model i generates the data. Define the log likelihood ratio

$$r_i \equiv \log \frac{L_{ii}}{L_{ij}},$$

where $j \neq i$ and i = A, B. When model *i* generates the data, r_i should be positive. Now consider the probabilities of two kinds of mistakes. First, assume that model A generates the data and calculate

$$p_A = \operatorname{Prob}\left(\operatorname{mistake}|A\right) = \operatorname{freq}\left(r_A \le 0\right).$$

 $^{^2}$ The two models (i.e., the two infinite-horizon stochastic processes) are locally absolutely continuous in the sense defined in Hansen, Sargent, Turuhambetova, and Williams (2001). The stochastic processes are not mutually absolutely continuous.

Details

Thus, p_A is the frequency of negative log likelihood ratios r_A when model A is true. Similarly, $p_B = \text{Prob}(\text{mistake}|B) = \text{freq}(r_B < 0)$ is the frequency of negative log likelihood ratios r_B when model B is true. Call the probability of a detection error

$$p\left(\theta\right) = \frac{1}{2}\left(p_A + p_B\right).$$

Here, θ is the robustness parameter used to generate a particular model B by taking the associated worst case perturbation of model A in light of a particular objective function for a decision maker. The following section shows in detail how to estimate the detection error probability by using simulations. In a given context, we propose to set $p(\theta)$ to a reasonable number, then invert $p(\theta)$ to find a plausible value of θ .

8.3. Details

We now describe how to compute detection error probabilities in some detail.

8.3.1. Likelihood ratio under the approximating model

Define w^A as the worst case shock assuming that the underlying data generating process is the approximating model, i.e., $w^A = \kappa x^A$ where x^A is generated under (8.1.1)). Define $\hat{A} = A_o + C\kappa$. Then we can express the innovation under the worst case model as:

$$\epsilon_{t+1} = (C'C)^{-1} C' \left(x_{t+1} - \hat{A}x_t \right), = \check{\epsilon}_{t+1} - \kappa x_t, = \check{\epsilon}_{t+1} - w_{t+1}^A.$$
(8.3.1)

The log likelihood function under the approximating model is

$$\log L_{AA} = -\frac{1}{T} \sum_{t=0}^{T-1} \{ \log \sqrt{2\pi} + \frac{1}{2} \left(\check{\epsilon}_{t+1} \cdot \check{\epsilon}_{t+1} \right) \}$$

The log likelihood function for the distorted model, given that the approximating model (8.1.1) is the data generating process, is

$$\log L_{AB} = -\frac{1}{T} \sum_{t=0}^{T-1} \{ \log \sqrt{2\pi} + \frac{1}{2} (\epsilon_{t+1} \cdot \epsilon_{t+1}) \},$$

$$= -\frac{1}{T} \sum_{t=0}^{T-1} \{ \log \sqrt{2\pi} + \frac{1}{2} (\check{\epsilon}_{t+1} - w_{t+1}^A)' (\check{\epsilon}_{t+1} - w_{t+1}^A) \}.$$
(8.3.2)

Hence, assuming that the approximating model is the data generating process, the likelihood ratio r_A is:

$$r_{A} \equiv \log L_{AA} - \log L_{AB},$$

= $\frac{1}{T} \sum_{t=0}^{T-1} \{ \frac{1}{2} w_{t+1}^{A} ' w_{t+1}^{A} - w_{t+1}^{A} ' \check{\epsilon}_{t+1} \}.$ (8.3.3)

The second term can be expected to vanish as $T \to \infty$, so the log likelihood ratio converges to the average value of the one-step measure of entropy $.5w_{t+1}^{A\prime}w_{t+1}^A$.

8.3.2. Likelihood ratio under the distorted model

Now suppose that the data generating process is the distorted model (8.1.2). The innovations in the approximating model are linked to those in the distorted model by $\check{\epsilon}_{t+1} = \epsilon_{t+1} + w^B_{t+1}$, where $w^B_{t+1} = \kappa x^B_t$ and x^B_t is generated under (8.1.2).

Assuming that the distorted model generates the data, the log likelihood function $\log L_{BB}$ for the distorted model is

$$\log L_{BB} = -\frac{1}{T} \sum_{t=0}^{T-1} \{ \log \sqrt{2\pi} + \frac{1}{2} \left(\epsilon_{t+1} \cdot \epsilon_{t+1} \right) \}.$$
(8.3.4)

The log likelihood function $\log L_{BA}$ for the approximating model, assuming that the distorted model (8.1.2) generates the data is,

$$\log L_{BA} = -\frac{1}{T} \sum_{t=0}^{T-1} \{ \log \sqrt{2\pi} + \frac{1}{2} \left(\check{\epsilon}_{t+1} \cdot \check{\epsilon}_{t+1} \right) \},$$

$$= -\frac{1}{T} \sum_{t=0}^{T-1} \{ \log \sqrt{2\pi} + \frac{1}{2} \left(\epsilon_{t+1} + w_{t+1}^B \right)' \left(\epsilon_{t+1} + w_{t+1}^B \right) \}.$$
(8.3.5)

Hence, the likelihood ratio r_B , assuming that the distorted model is the data generating process is

$$r_{B} \equiv \log L_{BB} - \log L_{BA},$$

= $\frac{1}{T} \sum_{t=0}^{T-1} \{ \frac{1}{2} w_{t+1}^{B'} w_{t+1}^{B} + w_{t+1}^{B'} \epsilon_{t+1} \}.$ (8.3.6)

As $T \to \infty$, this converges to the average value of one-period entropy $.5w_{t+1}^{B\prime}w_{t+1}^B$.

8.3.3. The detection error probability

Attach equal prior weights to model A and B. Then the detection error probability is

$$p(\theta) = \frac{1}{2}(p_A + p_B),$$
 (8.3.7)

where $p_i = \text{freq}(r_i \leq 0), i = A, B$. To compute $p(\theta)$, we simulate a large number of trajectories and calculate the empirical detection error probability.

8.3.4. Ball's model

We now illustrate the use of detection error probabilities to discipline the choice of θ in the context of the simple dynamic model that Ball (1999) designed to study alternative rules by which a monetary policy authority might set an interest rate.³ Ball's is a 'backward looking' macro model with the structure

$$y_t = -\beta r_{t-1} - \delta e_{t-1} + \epsilon_t \tag{8.3.8}$$

$$\pi_t = \pi_{t-1} + \alpha y_{t-1} - \gamma \left(e_{t-1} - e_{t-2} \right) + \eta_t \tag{8.3.9}$$

$$e_t = \theta r_t + \nu_t, \tag{8.3.10}$$

where y is the log of real output, r is the real interest rate, e is the log of the real exchange rate, π is the inflation rate, and ϵ , η , ν are serially uncorrelated and mutually orthogonal disturbances. As an objective, Ball assumed that the monetary authority wants to maximize

$$C = -E\left(\pi_t^2 + y_t^2\right).$$

 $^{^3\,}$ See Sargent (1999) for further discussion of Ball's model from the perspective of robust decision theory.

The government sets the interest rate r_t as a function of the current state at t, which Ball shows can be reduced to y_t, e_t .

Ball motivates (8.3.8) as an open-economy IS curve and (8.3.9) as an openeconomy Phillips curve; he uses (8.3.10) to capture effects of the interest rate on the exchange rate. Ball set the parameters $\gamma, \theta, \beta, \delta$ at the values .2, 2, .6, .2. Following Ball, we set the innovation shock standard deviations equal to 1, 1, $\sqrt{2}$.

To discipline the choice of the parameter expressing a preference for robustness, we calculated the detection error probabilities for distinguishing Ball's model from the worst-case models associated with various values of $\sigma \equiv -\theta^{-1}$. We calculated these taking Ball's parameter values as the approximating model and assuming that T = 142 observations are available, which corresponds to 35.5 years of annual data for Ball's quarterly model. Fig. 8.3.1 shows these detection error probabilities $p(\sigma)$ as a function of σ . Notice that the detection error probability is .5 for $\sigma = 0$, as it should be, because then the approximating model and the worst case model are identical. The detection error probability falls to .1 for $\sigma \approx -.085$. If we think that a reasonable preference for robustness is to want rules that work well for alternative models whose detection error probabilities are .1 or greater, then $\sigma = -.085$ is a reasonable choice of this parameter. In the next section, we'll compute a robust decision rule for Ball's model with $\sigma = -.085$ and compare its performance to the $\sigma = 0$ rule that expresses no preference for robustness.

8.3.5. Robustness in a simple macroeconomic model

We briefly illustrate how the detection error probabilities for Ball's model from Fig. 8.3.1 can be used to guide plausible the selection of defensible values of θ . We show a graph that quantifies the robustness attained by different settings of θ .

We use Ball's model to illustrate the robustness attained by alternative settings of the parameter θ . For Ball's model, we present Fig. 8.3.2 to show that while robust rules do less well when the approximating model actually generates the data, their performance deteriorates more slowly with departures of the data generating mechanism from the approximating model.

Fig. 8.3.2 plots the value $C = -E(\pi^2 + y^2)$ attained by three rules under the alternative data generating model associated with the worst case model for



Figure 8.3.1: Detection error probability (coordinate axis) as a function of $\sigma = -\theta^{-1}$ for Ball's model.

the value of σ on the ordinate axis.⁴ The rules are those for the three values $\sigma = 0, -.04, -.085$. Recall how the detection error probabilities computed above associate a value of $\theta = -0.085$ with a detection error probability of about .1. Notice how the robust rules (those computed with preference parameter $\sigma = -.04$ or -.085) have values that deteriorate at a lower rate with model misspecification (they are flatter). Notice that the rule for $\sigma = -.085$ does worse than the $\sigma = 0$ or $\sigma = -.04$ rules when $\sigma = 0$, but is more robust in deteriorating less when the model is misspecified.

⁴ Following the risk-sensitive control literature, we transform θ into the risk-sensitivity parameter $\sigma \equiv -\theta^{-1}$.



Figure 8.3.2: Value of $C = -E(\pi^2 + y^2)$ for three decision rules when the data are generated by the worst-case model associated with the value of σ on the horizontal axis: $\sigma = 0$ rule (solid line), $\sigma = -.04$ rule (dashed-dotted line), $\sigma =$ -.085 (dashed) line.

Chapter 9 A permanent income model

9.1. Introduction

Because economists have so much intuition about it, the permanent income model is a good laboratory for exploring the consequences of fears about model misspecification. A decision maker who distrusts his specification of his endowment process engages in a kind of precautionary savings that comes from his worst-case slanting of the probability law for the endowment process.¹ We use the Stackelberg multiplier game of chapter 6 to help us interpret how this probability slanting manifests itself in the permanent income model.

The permanent income model is also a good vehicle for gathering intuitions from the frequency domain approach of chapter 7. A permanent income consumer is patient enough to smooth high frequency fluctuations in income. But he is not be patient enough to smooth low frequency (i.e., very persistent) income fluctuations. Recognizing that low frequency income fluctuations cause the consumer the most trouble, the minimizing agent makes the worst case shocks persistent, an outcome that informs the consumer that his decision rule is most fragile with respect to low frequency misspecifications of the income process. The robust permanent income consumer responds to those more persistent worst case shocks by saving more than he would if he had no doubts about his endowment process. He engages in a form of precautionary savings that prevails even when he has quadratic preferences, which distinguishes it from the ordinary form of precautionary savings that emerges only with preferences that have convex marginal utilities.

We apply the label 'precautionary' because the effect increases with the volatility of innovations to endowments under the consumer's approximating model and because it also depends on the parameter θ that indexes his concern

- 221 -

¹ This context-specific slanting corresponds to that mentioned by Fellner in the passage cited on page 32 of chapter 1.

about robustness. Our model of precautionary savings exhibits the usual symptom that it modifies the certainty equivalence present in the linear-quadratic permanent income model. However, our model keeps the marginal propensity to save out of financial wealth equal to that out of human wealth, in contrast to models like those of Cabellero (XXXX) and BLANK (XXXX), where precautionary saving makes the marginal propensity to save out of human wealth exceed that out of financial wealth.²

To explore these issues, this chapter uses an equilibrium version of a permanent income model that Hansen, Sargent, and Tallarini (1999) (HST) estimated for U.S. consumption and investment data.³ We restate (and extend in appendix B) an observational equivalence result of HST, who showed that a concern about robustness increases saving just as would increasing the discount factor. Therefore, there exists an appropriate alteration of the discount factor that can offset the effect on the consumption and investment allocation of a change in the robustness parameter θ . HST thereby established that consumption and investment data alone are insufficient to identify both the robustness parameter θ and the subjective discount factor β .⁴ We use the Stackelberg multiplier game from chapter 6 to shed more light on the observational equivalence proposition and the impact on decision rules of distortions in the conditional expectations under the worst case model. We also state another observational equivalence result for a new baseline model and use it to show that activating a concern about robustness still equalizes the marginal propensities to save out of human and nonhuman wealth.

In addition, this chapter illustrates how the detection error probabilities of chapter 8 can discipline plausible choices of θ and provides some numerical examples of how much robustness can be achieved by rules designed with various settings of θ . In chapter 11, we describe how to decentralize the allocation chosen by the planner in the economy of this chapter. Then in chapter 12, we

 $^{^2}$ See Neng Wang (XXXX) for a treatment of how precautionary saving without robustness separates the marginal propensities to consumer out of financial and non-financial wealth.

³ Hall (1978), Campbell (1987), Heaton (1993), and Hansen, Roberds, and Sargent (1991) have applied versions of this model to aggregate U.S. time series data on consumption and investment.

⁴ Despite their failure to affect the consumption allocation, HST showed that such variations in (σ, β) do affect the relevant stochastic discount factor and therefore the valuation of risky assets. We shall take up asset pricing implications of the robust permanent income model in chapter 12.

use that decentralized economy as a laboratory for studying ways to represent the effects on asset prices of a concern about robustness.

9.2. A robust permanent income theory

HST's model features a planner with preferences over consumption streams $\{c_t\}_{t=0}^{\infty}$, intermediated through service streams $\{s_t\}_{t=0}^{\infty}$.⁵ Let b be a preference shifter in the form of a utility bliss point. The Bellman equation for the robust planner is

$$-x'Px - p = \sup_{c} \inf_{w} \left\{ -(s-b)^{2} + \beta \left(\theta w'w - Ey'Py - p \right) \right\}$$
(9.2.1) ["income1"]

where the maximization is subject to

$$s = (1 + \lambda) c - \lambda h$$
 (9.2.2*a*) ["income1a;a"]

$$h^* = \delta_h h + (1 - \delta_h) c \qquad (9.2.2b) \quad ["income1a;b"]$$

$$k^* = \delta_k k + i$$
 (9.2.2c) ["income1a;c "]

$$c+i=\gamma k+d$$
 (9.2.2d) ["income1a;d "]

$$\begin{bmatrix} d \\ b \end{bmatrix} = Uz_t \tag{9.2.2d} \quad ["income1a;d"]$$

$$z^* = A_{22}z + C_2 (\epsilon + w)$$
(9.2.2e) ["income1a;e"]
$$x = [h \ k \ z].$$
(9.2.2f) ["income1a;f"]

$$c = \begin{bmatrix} h & k & z \end{bmatrix}.$$
 (9.2.2f) ["income1a;f"]

Here * denotes next period's value, E is the expectation operator, c is consumption, s denotes a scalar service measure, and the law of motion mapping this period's state x into next period's state y will be defined below. As usual, the penalty parameter $\theta > 0$ governs concern about robustness to misspecification of the endowment process d and the preference shock process b embedded in (9.2.2d) and (9.2.2e). HST assumed that the eigenvalues of A_{22} are bounded in modulus by unity. We transform θ to the risk-sensitivity parameter $\sigma = -\theta^{-1}$. In (9.2.1), a scalar household service s_t is produced by the scalar consumption c_t via the household technology (9.2.2a) and (9.2.2b), namely,

$$s_t = (1 + \lambda) c_t - \lambda h_{t-1}$$
 (9.2.3*a*) ["income3;a"]

$$h_t = \delta_h h_{t-1} + (1 - \delta_h) c_t \qquad (9.2.3b) \quad [\texttt{"income3;b"}]$$

 $^{^5}$ The model fits within the framework described in chapter 10. See page 256 for an additional stability condition that must be imposed.

where $\lambda > 0$ and $\delta_h \in (0, 1)$. System (9.2.3) accommodates habit persistence or durability as in Ryder and Heal (1973), Becker and Murphy (1988), Sundaresan (1989), Constantinides (1990) and Heaton (1993). By construction, h_t is a geometric weighted average of current and past consumption. Setting $\lambda > 0$ induces intertemporal complementarities. Consumption services depend positively on current consumption, but negatively on a weighted average of past consumption, a reflection of 'habit persistence'.

There is a linear production technology (9.2.2d) where the capital stock k^* at the end of period t evolves according to (9.2.2e), where i_t is time t gross investment, and $\{d_t\}$ is an exogenously specified endowment process. The parameter γ is the (constant) marginal product of capital, and δ_k is the depreciation factor for capital. HST specified a bivariate ('two-factor') stochastic endowment process: $d_t = \mu_d + d_t^* + \hat{d}_t$.⁶ They assumed that the two endowment processes are orthogonal and that both obey second order autoregressions:

$$(1 - \phi_1 L) (1 - \phi_2 L) d_t^* = c_{d^*} \left(\epsilon_t^{d^*} + w_t^{d^*} \right)$$
$$(1 - \alpha_1 L) (1 - \alpha_2 L) \hat{d}_t = c_{\hat{d}} \left(\epsilon_t^{\hat{d}} + w_t^{\hat{d}} \right)$$

where the vector ϵ_t is i.i.d. Gaussian with mean zero and identity covariance matrix, and $w_t^{d^*}, w_t^{\hat{d}}$ are distortions to the means of $\epsilon_t^{d^*}, \epsilon_t^{\hat{d}}$. HST estimated values of the ϕ_j 's and α_j 's that imply that the d_t^* process is more persistent than the \hat{d}_t process, as we see below.

Solving the capital evolution equation for investment and substituting into the linear production technology gives

$$c_t + k_t = Rk_{t-1} + d_t$$
 (9.2.4) ["income4"]

where

$$R \equiv \delta_k + \gamma$$

which is the physical gross return on capital, taking into account that capital depreciates over time.⁷

⁶ For two observed time series (c_t, i_t) , HST's econometric specification needed at least two shock processes to avoid 'stochastic singularity'.

 $^{^{7}}$ For HST's decentralized economy, R coincided with the gross return on a risk free asset.

The state vector can be taken to be $x'_t = \begin{bmatrix} h_{t-1} & k_{t-1} & d_{t-1} & 1 & d_t & d^*_t & d^*_{t-1} \end{bmatrix}'$ (see Hansen, Sargent, and Wang (2002)). There is a set of state transition equations indexed by a $\{w_{t+1}\}$ process:

$$x_{t+1} = Ax_t + Bu_t + C(w_{t+1} + \epsilon_{t+1})$$
(9.2.5) ["law0 "]

where $u_t = c_t$ and $w'_{t+1} = \begin{bmatrix} w_{t+1}^d & w_{t+1}^d \end{bmatrix}'$ is the distortion to the conditional mean of ϵ_{t+1} . Let J_t be the sigma algebra induced by $\{x_0, \epsilon_s, 0 \le s \le t\}$. We impose that the components of the solution for $\{c_t, h_t, k_t\}$ belong to L_0^2 , the space of stochastic processes $\{y_t\}$ defined as:

$$L_0^2 = \{y : y_t \text{ is in } J_t \text{ for } t = 0, 1, \dots \text{ and } E \sum_{t=0}^{\infty} R^{-t} (y_t)^2 \mid J_0 < +\infty\}.$$

Given x_0 , the planner chooses a process $\{c_t, k_t\}$ with components in L_0^2 to solve the Bellman equation (9.2.1) subject to versions of (9.2.3), (9.2.4).⁸ Soon we'll discuss HST's parameter values and some properties of their numerical solution. But first we show that in terms of its effects on consumption and investment, more concern about robustness works just like an increase in the discount factor.⁹

⁸ We can convert this problem into a special case of the control problem posed in chapter 6 as follows. Form a composite state vector x_t as described above, and let the control be given by $s_t - b_t$. Solve (9.2.3*a*) for c_t as a function of $s_t - b_t$, b_t and h_{t-1} and substitute into equations (9.2.3*b*) and (9.2.4). Stack the resulting two equations along with the state evolution equation for z_t to form the evolution equation for x_{t+1} .

⁹ However, in chapter 12, we show that (β, σ) pairs that are observationally equivalent for consumption and investment nevertheless imply different prices for risky assets.

9.3. Solution when $\sigma = 0$

We apply results from chapter 6 to show that the robust decision rule for $\sigma < 0$ also solves a $\sigma = 0$ version of the model in which the maximizing agent in (9.2.1) replaces the approximating model with a particular distorted model for $[d_t, b_t]$. We shall eventually use that insight to study the identification of σ and β . To begin, this section solves the $\sigma = 0$ model.

9.3.1. The $\sigma = 0$ benchmark case

This subsection computes a solution of the planning problem in the $\sigma = 0$ case. Though we shall soon focus on the case when $\beta R = 1$, we also want the solution when $\beta R \neq 1$. Thus, for now we allow $\beta R \neq 1$. When $\sigma = 0$, the decision maker's objective reduces to

$$E_0 \sum_{t=0}^{\infty} \beta^t \{ -(s_t - b_t)^2 \}.$$
 (9.3.1) ["income5"]

Formulate the planning problem as a Lagrangian by putting random Lagrange multiplier processes $2\beta^t \mu_{st}$ on (9.2.3*a*), $2\beta^t \mu_{ht}$ on (9.2.3*b*), and $2\beta^t \mu_{ct}$ on (9.2.4). First-order necessary conditions are

$$\mu_{st} = b_t - s_t \tag{9.3.2a} \quad ["income6;a"]$$

$$\mu_{ct} = (1 + \lambda) \,\mu_{st} + (1 - \delta_h) \,\mu_{ht} \tag{9.3.2b} \quad \texttt{["income6;b"]}$$

$$\mu_{ht} = \beta E_t \left[\delta_h \mu_{ht+1} - \lambda \mu_{st+1} \right]$$
(9.3.2c) ["income6;c"]

$$\mu_{ct} = \beta R E_t \mu_{ct+1} \tag{9.3.2d} \quad \texttt{["income6;d"]}$$

and also (9.2.3), (9.2.4). Equation (9.3.2d) implies that $E_t \mu_{ct+1} = (\beta R)^{-1} \mu_{ct}$. Then (9.3.2b) and (9.3.2c) solved forward imply that μ_{st}, μ_{ht} must satisfy $E_t \mu_{st+1} = (\beta R)^{-1} \mu_{st}$ and $E_t \mu_{ht+1} = (\beta R)^{-1} \mu_{ht}$. Therefore, μ_{st} has the representation

$$\mu_{st} = \left(\beta R\right)^{-1} \mu_{st-1} + \nu' \epsilon_t \qquad (9.3.3) \quad ["martingale"]$$

for some vector ν . The endogenous volatility vector ν will play an important role below, and we shall soon tell how to compute it.

Use (9.3.2a) to write $s_t = b_t - \mu_{st}$, substitute this into the household technology (9.2.3), and rearrange to get the system

$$c_t = \frac{1}{1+\lambda} \left(b_t - \mu_{st} \right) + \frac{\lambda}{1+\lambda} h_{t-1}$$
(9.3.4a) ["income7;a"]

$$h_t = \tilde{\delta}_h h_{t-1} + \left(1 - \tilde{\delta}_h\right) \left(b_t - \mu_{st}\right)$$
(9.3.4b) ["income7;b"]

where $\tilde{\delta}_h = \frac{\delta_h + \lambda}{1 + \lambda}$. Equation (9.3.4b) can be used to compute

$$E_t \sum_{j=0}^{\infty} R^{-j} h_{t+j-1} = \left(1 - R^{-1} \tilde{\delta}_h\right)^{-1} h_{t-1} + \frac{R^{-1} \left(1 - \tilde{\delta}_h\right)}{\left(1 - R^{-1} \tilde{\delta}_h\right)} E_t \sum_{j=0}^{\infty} R^{-j} \left(b_{t+j} - \mu_{st+j}\right).$$

$$(9.3.5) \quad ["income8"]$$

For the purpose of solving the first-order conditions (9.3.2), (9.2.3), (9.2.4) subject to the side condition that $\{c_t, k_t\} \in L_0^2$, treat the technology (9.2.4) as a difference equation in $\{k_t\}$, solve forward, and take conditional expectations on both sides to get

$$k_{t-1} = \sum_{j=0}^{\infty} R^{-(j+1)} E_t \left(c_{t+j} - d_{t+j} \right).$$
(9.3.6) ["income9"]

Use (9.3.4*a*) to eliminate $\{c_{t+j}\}$ from (9.3.6), then use (9.3.3) and (9.3.5). Solve the resulting system for μ_{st} to get

$$\mu_{st} = \Psi_1 k_{t-1} + \Psi_2 h_{t-1} + \Psi_3 \sum_{j=0}^{\infty} R^{-j} E_t b_{t+j} + \Psi_4 \sum_{j=0}^{\infty} R^{-j} E_t d_{t+j}, \quad (9.3.7) \quad ["income10"]$$

where

$$\begin{split} \Psi_{1} &= -(1+\lambda) R \left(1 - R^{-2} \beta^{-1}\right) \left[\frac{1 - R^{-1} \tilde{\delta}_{h}}{1 - R^{-1} \tilde{\delta}_{h} + \lambda \left(1 - \tilde{\delta}_{h}\right)} \right] \\ \Psi_{2} &= \frac{\lambda \left(1 - R^{-2} \beta^{-1}\right)}{1 - R^{-1} \tilde{\delta}_{h} + \lambda \left(1 - \tilde{\delta}_{h}\right)} \end{split}$$
(9.3.8) ["income100 "]
$$\Psi_{3} &= \left(1 - R^{-2} \beta^{-1}\right) \\ \Psi_{4} &= R^{-1} \Psi_{1}. \end{split}$$

Equations (9.3.7), (9.3.4), and (9.2.4) represent the solution of the planning problem when $\sigma = 0.10$

¹⁰ When $\beta R = 1$, (9.3.7) makes μ_{st} depend on a geometric average of current and future values of b_t . Therefore, the optimal consumption service process and optimal consumption both depend on the difference between b_t and a geometric average of current and expected future values of b. So there is no 'level effect' of the preference shock on the optimal decision rules for consumption and investment. However, the level of b_t will affect equilibrium asset prices.

To compute ν in (9.3.3), it is useful to notice that formula (9.3.7) can be rewritten as

$$\mu_{st} = (\beta R)^{-1} \mu_{st-1} + \Phi_3 \sum_{t=0}^{\infty} R^{-j} (E_t b_{t+j} - E_{t-1} b_{t+j})$$

$$+ \Phi_4 \sum_{t=0}^{\infty} R^{-j} (E_t d_{t+j} - E_{t-1} d_{t+j})$$
(9.3.9) ["murepresentation "]

where

$$\mu_{st-1} = \Phi_1 k_{t-1} + \Phi_2 h_{t-1} + \Phi_3 \sum_{t=0}^{\infty} R^{-j} E_{t-1} b_{t+j} + \Phi_4 \sum_{t=0}^{\infty} R^{-j} E_{t-1} d_{t+j}$$

The terms $\Phi_3 \sum_{t=0}^{\infty} R^{-j} (E_t b_{t+j} - E_{t-1} b_{t+j})$ and $\Phi_4 \sum_{t=0}^{\infty} R^{-j} (E_t d_{t+j} - E_{t-1} d_{t+j})$ are scalars Ψ_3 and Ψ_4 times the innovations at t in the present values of b_t and d_t , respectively. Let the moving average representations for b_t and d_t be

$$b_{t} = \zeta_{b}\left(L\right)\epsilon_{t} \tag{9.3.10} \quad \texttt{["mab "]}$$

$$d_t = \zeta_d \left(L \right) \epsilon_t, \tag{9.3.11} \quad \texttt{["mad "]}$$

where from (9.2.2e) $\zeta_b(L) = U_b(I - A_{22}L)^{-1}C_2$ and $\zeta_d(L) = U_d(I - A_{22}L)^{-1}C_2$. By applying a formula of Hansen and Sargent (1980XXX), it is easy to show that the innovations in the present values of b_t and d_t , respectively, equal the present values of the moving average coefficients in these moving average representations.¹¹ Therefore, representation (9.3.9) can be rewritten as

$$\mu_{st} = (\beta R)^{-1} \,\mu_{st-1} + \left[\Psi_3 \zeta_b \left(R^{-1} \right) + \Psi_4 \zeta_d \left(R^{-1} \right) \right] \epsilon_t. \tag{9.3.12} \quad \texttt{["murep2"]}$$

Comparing this with (9.3.3), we see that

$$\nu = \Psi_3 \zeta_b \left(R^{-1} \right) + \Psi_4 \zeta_d \left(R^{-1} \right). \tag{9.3.13} \quad ["hsoffset1"]$$

An equivalent way to compute ν is to note that formula (9.3.7) for μ_{st} can be represented in matrix notation as

$$\mu_{st} = M_s x_t \tag{9.3.14} \quad ["muslaw "]$$

$$x_t = A_o x_{t-1} + C\epsilon_t$$
 (9.3.15) ["xlaw"]

 $^{^{11}}$ The present value of the moving average coefficients plays an important role in linear quadratic permanent income models. See Flavin (XXX), Campbell (XXX), and Hansen, Roberds, and Sargent (xXXX).

where x_t is the state vector k_{t-1}, h_{t-1}, z_t , where $z_t = \begin{bmatrix} d_{t-1} & 1 & d_t & d_t^* & d_{t-1}^* \end{bmatrix}'$ the matrix M_s is determined by equations (9.3.7) and the laws of motion for b_t, d_t , and A_o, C tell the law of motion for the entire state under the optimal rule for c_t .¹² It follows that $\mu_{st} = M_s A_o x_{t-1} + M_s C \epsilon_t$, which must agree with (9.3.3), so that $\mu_{s,t-1} \equiv M_s A_o x_{t-1}$ and

$$\nu' \equiv M_s C$$
 (9.3.16) ["hsoffset2 . "]

229

The scalar $\alpha = \sqrt{\nu'\nu}$ plays an important role in the argument below. It obeys

$$\alpha = \sqrt{M_s C C' M'_s}. \tag{9.3.17} ["MsC"]$$

In the widely studied special case that $\lambda = \delta_h = 0$, so that $s_t = c_t$ and $\mu_{st} = b_t - c_t$, (9.3.7), (9.3.8) imply that the marginal propensity to consume out of "non-human wealth" defined as Rk_{t-1} and the marginal propensity to consume out of "human wealth" defined as $\sum_{j=0}^{\infty} R^{-j} E_t d_{t+j}$ both equal $-\Psi_1$. It is a well known feature of the linear-quadratic model that these marginal propensities to consume are equal. Notice that human wealth is formed by discounting expected future endowments at the risk-free rate.

9.3.2. Observational equivalence (for quantities) of $\sigma = 0$ and $\sigma \neq 0$

In the $\sigma = 0$ case, HST followed Hall (1978) and imposed that $\beta R = 1$. HST then showed that for fixed values of all other parameters, there is a set of (β, σ) pairs that leave the consumption-investment plan unaltered. In particular, if as we vary σ we also vary β according to ¹³

$$\hat{\beta}\left(\sigma\right) = \frac{1}{R} + \frac{\sigma\alpha^2}{R-1}, \qquad (9.3.18) \quad ["obseq"]$$

then we leave unaltered the decision rules for (c_t, i_t) . Here $\alpha^2 = \nu'\nu$, where ν as defined in (9.3.13) is a vector in the following martingale representation for the marginal utility of services μ_{st} that prevails as a special case of (9.3.3) when $\sigma = 0$ and $\beta R = 1$:

$$\mu_{st} = \mu_{st-1} + \nu' \epsilon_t.$$

¹² Here C is the matrix that appears in (9.2.5) above. See Hansen and Sargent (20XX) for fast ways to compute A_o, M_s, C for a class of models that includes that of this chapter. ¹³ See footnote 22.

(Also see equation (9.3.12)). This section explains how HST constructed the locus identified by (9.3.18).

9.3.3. Observational equivalence: intuition

Here is the basic idea underlying the observational equivalence proposition. A single factor μ_{st} effectively summarizes all of the endogenous state variables in the model. When $\beta R = 1$ and $\sigma = 0$, it has the law of motion

$$\mu_{st} = \mu_{st-1} + \nu' \epsilon_t$$

which can also be represented as

$$\mu_{st} = \mu_{st-1} + \alpha \tilde{\epsilon}_t \tag{9.3.19} \quad ["reversee1 "]$$

where $\tilde{\epsilon}_t$ is a scalar i.i.d. process with zero mean and unit variance and α verifies $\alpha \tilde{\epsilon}_t = \nu' \epsilon_t$. We generate our observational equivalence result by reverse engineering. We activate a concern about robustness by setting $\sigma < 0$, but insist that (9.3.19) continue to be the approximating model for μ_{st} . In this way, we freeze the (c_t, i_t) allocation. For $\sigma < 0$, the worst-case model for μ_{st} is evidently then

$$\mu_{st} = \mu_{st-1} + \alpha \left(\tilde{\epsilon}_t + \tilde{w}_t \right) \tag{9.3.20} \quad ["reversee2"]$$

or

$$\mu_{st} = (1 + \alpha K(\sigma)) \mu_{st-1} + \alpha \tilde{\epsilon}_t \qquad (9.3.21) \quad ["reversee3"]$$

where $w_t = K(\sigma)\mu_{st-1}$. Now with a concern about robustness, the decision maker's choices conform to the Euler equation $\hat{E}_t\mu_{st+1} = (\hat{\beta}R)^{-1}\mu_{st}$ where \hat{E}_t is evaluated with respect to the distorted model (9.3.21) and $\hat{\beta}$ is a new value for β . We want this distorted model to be associated with the approximating model (9.3.19). But according to (9.3.21), if the approximating model is to be (9.3.19), then $\hat{E}_t\mu_{st+1} = (1 + K(\sigma)\alpha)\mu_{st}$. Thus, we want to find a replacement $\hat{\beta}$ for β that enables us to verify $(\hat{\beta}R)^{-1} = (1 + \alpha K(\sigma))$, where $K(\sigma)$ solves the minimization problem that gives rise to the worst case shock. In effect, we want to solve $(\hat{\beta}R)^{-1} = 1 + \alpha K(\sigma)$ for $\hat{\beta}$ as a function of σ . The formal proof of observational equivalence shows that a $\hat{\beta}$ that satisfies (9.3.18) does the job.

9.3.4. Observational equivalence: formal argument

Following HST, we begin by assuming that $\beta R = 1$ when $\sigma = 0$. We state

Theorem 9.3.1. (Observational Equivalence, I) Fix all parameters, including R, except (β, σ) . Suppose $\beta R = 1$ when $\sigma = 0$. There exists a $\underline{\sigma} < 0$ such that for any $\sigma \in (\underline{\sigma}, 0)$, the optimal consumption-investment plan for $(\beta, 0)$ is also chosen by a robust decision maker when parameter values are $(\hat{\beta}(\sigma), \sigma)$ and where $\hat{\beta}(\sigma) < \beta$ satisfies (9.3.18).

Proof. The proof of the proposition is constructive. Begin with a solution $\{\bar{s}_t, \bar{c}_t, \bar{k}_t, \bar{h}_t\}$ for a benchmark $\sigma = 0, \beta R = 1$ economy, then form a comparison economy with a $\sigma \in [\underline{\sigma}, 0]$, where $\underline{\sigma}$ is the lowest value for which the solution of (9.3.25) reported below is real. The comparison economy fixes all parameters except (σ, β) at their values for the benchmark economy. We then construct a discount factor $\hat{\beta} < \beta$ for which $\{\bar{s}_t, \bar{c}_t, \bar{k}_t, \bar{h}_t\}$ is also the allocation for the $\sigma < 0$ economy.

When $\beta R = 1$, (9.3.3) becomes

$$\mu_{st} = \mu_{st-1} + \nu' \epsilon_t. \tag{9.3.22} ["neng3 "]$$

The optimality of the allocation under the original $(0,\beta)$ implies that (9.3.22)is satisfied, which in turn implies that $E_t\mu_{ct+1} = \mu_{ct}$ and (9.3.7) are satisfied where E_t is the expectation operator under the approximating model. We seek a new value $\sigma < 0$ and an associated value $\hat{\beta}(\sigma)$ for which: (1) (9.3.22) remains satisfied under the approximating model; (2) the robust decision maker chooses the $(\bar{\cdot})$ allocation, which requires that $\hat{\beta}R\hat{E}_t\mu_{ct+1} = \mu_{ct}$,¹⁴ where \hat{E} is the expectation with respect to the worst case model associated with $(\sigma, \hat{\beta})$ when the approximating model obeys (9.3.22). However, when the approximating model satisfies (9.3.22), the worst case model associated with $(\sigma, \hat{\beta})$ implies that $\hat{E}_t\mu_{ct+1} = \hat{\zeta}(\hat{\beta})\mu_{ct}$, where $\hat{\zeta} > 1$ can be found by solving the pure forecasting problem¹⁵ associated with law of motion (9.3.22), one-period return function $-\mu_{st}^2 = -(b_t - s_t)^2$, and discount factor $\hat{\beta}$. If the σ -robust decision maker is to choose a decision rule that sustains (9.3.22) under the approximating model, so that (1) and (2) both prevail, $\hat{\beta}$ must verify

$$\hat{eta}R\hat{\zeta}\left(\hat{eta}
ight)=1.$$
 (9.3.23) ["eulerdist"]

 $^{^{14}}$ This is the robust decision maker's Euler equation for capital.

 $^{^{15}}$ See page 172 for the definition of a pure forecasting problem.

To complete the argument, we compute $\hat{\zeta}(\hat{\beta})$ by solving a pure forecasting problem to find the distorted expectation operator \hat{E}_t . We use the recipe given in formulas (6.B.9) on page 169 and (6.B.25) and (6.B.26) on page 172. Taking (9.3.22) as given under the approximating model and noting that $\mu_{st}^2 = (b_t - s_t)^2$, the evil agent in the pure forecasting problem seeks to minimize $-\sum_{t=0}^{\infty} \beta^t (\mu_{st}^2 + \beta_{\sigma}^1 w_{t+1}^2)$ under the distorted law $\mu_{st} = \mu_{st-1} + \alpha w_t$, where $\alpha = \sqrt{\nu'\nu}$ (see (9.3.22)). Taking μ_s as the state, the evil agent's Bellman equation (6.B.26) is¹⁶

$$-P\mu_s^2 = -\mu_s^2 + \beta \min_{w} \left(-\frac{1}{\sigma} w^2 - P\left(\mu_s + \alpha w\right)^2 \right).$$
(9.3.24) ["pess1"]

The scalar P that solves (9.3.24) is

$$-P\left(\beta\right) = \frac{\beta - 1 + \sigma\alpha^2 + \sqrt{\left(\beta - 1 + \sigma\alpha^2\right)^2 + 4\sigma\alpha^2}}{-2\sigma\alpha^2}.$$
 (9.3.25) ["distortcons"]

Let $\hat{\zeta} = A + CK = 1 + \alpha K$, where $w = K\mu_s$ is the formula for the worst case shock and A + CK is the state transition matrix for the distorted law of motion in chapter 6. Applying formula (6.B.20) for K in chapter 6 to the current problem gives

$$\hat{E}_t \mu_{st+1} = \hat{\zeta} \mu_{st}$$
 (9.3.26) ["income11"]

where

$$\hat{\zeta} = \hat{\zeta}\left(\beta\right) = 1 + \frac{\sigma\alpha^2 P\left(\beta\right)}{1 - \sigma\alpha^2 P\left(\beta\right)} = \frac{1}{1 - \sigma\alpha^2 P\left(\beta\right)}.$$
(9.3.27) ["distort2"]

Hansen, Sargent, and Wang (2002) solve (9.3.23), (9.3.25), and (9.3.27) to obtain

$$\hat{\beta}\left(\sigma\right) = \frac{1}{R} + \frac{\sigma\alpha^2}{R-1}. \tag{9.3.28} \quad ["obseqtheorem"]$$

For $\sigma \in [\underline{\sigma}, 0]$, equation (9.3.28) defines a locus of $(\sigma, \hat{\beta})$'s, each point of which is observationally equivalent to $(0, \beta)$ for observations on (c_t, k_t) because each supports the benchmark $(\sigma = 0)$ allocation.

This proposition means that with the appropriate adjustments in β given by $\hat{\beta}(\sigma)$, the robust decision maker chooses precisely the same quantities $\{c_t, k_t\}$ as

¹⁶ We exploit certainty equivalence and ignore the stochastic parts of the Bellman equation and the law of motion for μ_s .

a decision maker without a concern for robustness. Thus, as far as these quantity observations are concerned, the robust ($\sigma < 0$) version of the permanent income model is observationally equivalent to the benchmark ($\sigma = 0$) version.¹⁷

9.3.5. Precautionary savings interpretation

The consumer's concern about model misspecification activates a particular kind of precautionary savings motive that underlies our observational equivalence proposition. A concern about robustness inspires the consumer to save *more*. Decreasing his discount factor induces the consumer to save *less*. The observational equivalence proposition asserts that these two effects can be forced to offset each other.

The following experiment highlights the precautionary motive for savings. Take the base model with $\sigma = 0$ used in our proof of Theorem 9.3.1. Then activate a concern about robustness by setting $\sigma < 0$, but offset its effect on consumption by setting β equal to $\hat{\beta}(\sigma)$. Notice that $\hat{\beta}(\sigma)$ depends on the volatility parameter α . Consider a $(\sigma, \hat{\beta}(\sigma))$ pair corresponding to a given $\alpha > 0$. The innovation volatility associated with a positive α means that future endowments are forecast with error. If future endowments and preference shifters *could* be forecast perfectly, then at the value $\beta = \hat{\beta}(\sigma)$, the consumer would choose to make his capital stock, and therefore also his consumption, drift downward because discounting is large relative to the marginal productivity of capital. Investment would be sufficiently unattractive that the optimal linear

¹⁷ The asset pricing theory developed by HST which is encoded in (9.3.23) implies that the price of a sure claim on consumption one period ahead is R^{-1} for all t and for all $(\sigma, \hat{\beta})$ in the locus (9.3.18). Therefore, these different parameter pairs are also observationally equivalent with respect to the risk-free rate. In this model, the technology (9.2.4) ties down the risk-free rate. For a version of the model with quadratic costs of adjusting capital, the risk-free rate comes to depend on σ , even though the observations on quantities are nearly independent of σ . See Hansen and Sargent (1996).

rule would eventually send both consumption and capital below zero.¹⁸, ¹⁹ However, when randomness is activated (i.e., the innovation variances are positive), this downward drift is arrested or even completely offset, as it is in our observation equivalence proposition. Thus our robust control interpretation of the permanent income decision rule delivers a form of precautionary savings.

The precautionary savings coming from a concern about robustness differs in structure from another, perhaps more familiar, kind of precautionary savings that has attracted much attention in the macroeconomics literature and that emerges when a positive *variance* of the innovations to the endowment process interacts with a convex derivative of the marginal utility of consumption.²⁰ In contrast, the precautionary savings induced by a concern about robustness emerges because the consumer wants to protect himself against mistakes in specifying conditional *means* of shocks to the endowment. Thus, a concern for robustness inspires precautionary savings because of the way fears of misspecification are expressed in conditional *first* moments of shocks. This type of precautionary saving does not require that the marginal utility of consumption be convex and occurs even in models with quadratic preferences.

A concern about robustness affects consumption by slanting probabilities in the way described by Fellner in the passage cited on page 32 of this book. The household saves more for a given β because it makes pessimistic forecasts of future endowments. Precisely how pessimism manifests itself depends on the detailed structure of the permanent income model and the temporal properties of the endowment process, as we shall discuss in the next section.

¹⁸ Introducing nonnegativity constraints in capital and/or consumption would induce nonlinearities into the consumption and savings rules, especially near zero capital. But investment would remain unattractive in the presence of those constraints for experiments like the one we are describing here. See Deaton (1991) for a survey and quantitative assessment of consumption models with binding borrowing constraints.

¹⁹ As emphasized by Carroll (1992), even when the discount factor is small relative to the interest rate, precautionary savings can emerge when there is a severe utility cost for zero consumption. Such a utility cost is absent in our formulation.

²⁰ Take the Euler equation $E_t\beta Ru'(c_{t+1}) = u'(c_t)$ and assume that $\beta R = 1$ so that $E_t u'(c_{t+1}) = u'(c_t)$. If u' is a convex function, then applying Jensen's inequality implies $E_t c_{t+1} > c_t$, so that consumption is expected to grow when the conditional distribution of c_{t+1} is not concentrated at a point. Such consumption growth reflects precautionary savings. See Ljungqvist and Sargent (2004XXXX, chapter 16) for an analysis of these precautionary savings models.

9.4. Observational equivalence and distorted expectations

In this section, we use insights from the Stackelberg multiplier game on page 155 to interpret Theorem 9.3.1. In the Stackelberg multiplier game, decisions for the maximizing player can be computed by solving his Euler equations using a particular distorted law of motion to form conditional expectations of the shocks.²¹

In the benchmark $\sigma = 0, \beta R = 1$ case that is contemplated in Theorem 9.3.1, the solution of the planning problem is determined by equations (9.3.4), (9.2.4), and (9.3.7) where the Ψ_j 's satisfy (9.3.8) with $\beta R = 1$. For a $\sigma \in [\underline{\sigma}, 0)$ and a $\beta = \hat{\beta}(\sigma)$, the decision rule for the robust planner is characterized by equations (9.3.4), (9.2.4), and the following modified version of (9.3.7):

$$\mu_{st} = \hat{\Psi}_1 k_{t-1} + \hat{\Psi}_2 h_{t-1} + \hat{\Psi}_3 \sum_{j=0}^{\infty} R^{-j} \hat{E}_t b_{t+j} + \hat{\Psi}_4 \sum_{j=0}^{\infty} R^{-j} \hat{E}_t d_{t+j}, \quad (9.4.1) \quad \texttt{["income101"]}$$

where $\hat{\Psi}_j$ are determined by (9.3.8) with $\beta = \hat{\beta}(\sigma)$; and \hat{E}_t is the conditional expectation operator with respect to the distorted law of motion for the state x_t . The observational equivalence Theorem 9.3.1 implies that (9.4.1) and (9.3.7) are identical solutions for μ_{st} . By eliminating the terms in expected future values, the solutions (9.3.7) and (9.4.1) can also be expressed as $\mu_{st} = M_s x_t$ and $\mu_{st} = \hat{M}_s x_t$. Observational equivalence requires that $M_s = \hat{M}_s$. This requires that the $\hat{\Psi}_j$'s and \hat{E} mutually adjust to keep M_s fixed.²²

To expand on this point, consider the special case that $\lambda = \delta_h = 0$, so that we need not retain h_{t-1} as a state variable. Also, assume for simplicity that $b_t = b$, so that the preference shock is constant. Shutting down the volatility of b prevents distortions in it from affecting the robust decision rule. Then

 $^{^{21}}$ While the timing protocol for the Stackelberg multiplier game differs from the Markov perfect timing embedded in game (9.2.1), chapter 6 showed that identical equilibrium outcomes and recursive representations of equilibria prevail under the two timing protocols.

²² Note from formula (**@Eq.hsoffset2@**) that M_s determines α , a key parameter defining the observational equivalence locus (9.3.18). Thus, because M_s remains fixed, so does α so long as ($\hat{\beta}, \sigma$) respect (9.3.18).

equating the right sides of (9.3.7) and (9.4.1) gives

$$0 = \left(\Psi_4 - \hat{\Psi}_4\right) Rk_{t-1} + \left(\Psi_3 - \hat{\Psi}_3\right) \left(1 - R^{-1}\right)^{-1} b + \Psi_4 \sum_{j=0}^{\infty} R^{-j} E_t d_{t+j} - \hat{\Psi}_4 \sum_{j=0}^{\infty} R^{-j} \hat{E}_t d_{t+j}$$
(9.4.2) ["newequiv"]

where Ψ_j without hats denotes values of Ψ_j that satisfy (9.3.8) and those with hats satisfy (9.3.8) evaluated at $\beta = \hat{\beta}(\sigma)$. Equation (9.4.2) shows how the observational equivalence result asserts offsetting alterations in the coefficients Ψ_j and the distorted expectations operator \hat{E}_t used to form the expected sum of discounted future endowments that defines human wealth.²³

The distorted expectations operator is to be interpreted in terms of the recursive formulation of the maximizing player's problem in game 1 of chapter 6, the multiplier game game in sequences. (see section 6.5.1, pages 150–155). The Euler equation approach used to derive (9.3.7) or (9.4.1) presumes the following timing protocol. After the minimizing player has committed to an entire path for the w_{t+1} process, the maximizing agent faces the following recursive representation of the motion for the endowment and preference shocks:

$$\begin{aligned} X_{t+1} &= \left(A - BF\left(\sigma\right) + CK\left(\sigma\right)\right)X_t + C\tilde{\epsilon}_{t+1} & (9.4.3a) \quad ["\texttt{sys2;a"}] \\ \begin{bmatrix} b_t \\ d_t \end{bmatrix} &= SX_t & (9.4.3b) \quad ["\texttt{sys2;b"}] \end{aligned}$$

where $\tilde{\epsilon}_{t+1}$ is an i.i.d. shock identical in distribution to that of ϵ_{t+1} .²⁴ Because the minimizing player has committed himself to a stochastic process for $\{w_{t+1}\}$ that implies the recursive representation (9.4.3) of the endowment and preference shock processes, the maximizing player takes the X_t process as exogenous and uses the forecasting rule $\hat{E}_t X_{t+j} = (A - BF(\sigma) + CK(\sigma))^j X_t$ to form forecasts of (b_{t+j}, d_{t+j}) in (9.4.1). These forecasts, together with (9.4.1), (9.3.4), and (9.2.4) can be solved as in chapter 6 for a decision rule $c_t = -\mathcal{F}\begin{bmatrix} x_t \\ X_t \end{bmatrix}$. After computing the decision rule as a function of x_t, X_t , we equate $x_t = X_t$; that gives the maximizing agent's decision rule in the form $c_t = -Fx_t$.²⁵

 $^{^{23}}$ XXXXX We confirmed this in the program hst4.m in the subdirectory hst.

²⁴ In (9.4.3), X_t is used to attain a recursive representation of the distorted endowment and preference shock process and to keep it exogenous to the maximizer's decisions.

²⁵ The procedure of first optimizing, then setting $x_t = X_t$ to eliminate X_t is a common way of formulating rational expectations equilibria in macroeconomics, where it is sometimes called the 'Big K, little k' method.

9.4.1. Distorted endowment process

Fig. 9.4.1 and Fig. 9.4.2, illustrate the probability slanting that leads to precautionary savings. The figures assume HST's parameter values (see Appendix A)²⁶ and record impulse response functions for the total endowment d_t under the approximating model and a worst-case model associated with $\sigma = -.0001$, where β is adjusted according to (9.3.18) as required under our observational equivalence proposition to preserve the same decision rule $F(\sigma)$ for different σ 's.²⁷

For the approximating and the worst case model for $\sigma = -.0001$, the figures report the response of the total endowment d_t to innovations ϵ_t^* and $\hat{\epsilon}_t$ in the relatively permanent and transitory components of the endowment, d_t^*, \hat{d}_t , respectively. Under the distorted model, the impulse response functions diverge and the eigenvalue of $A - BF(\sigma) + CK(\sigma)$ that has maximum modulus increases from its value of unity under the approximating model to 1.0016.

The distorted endowment processes respond to innovations with more persistence than they do under the approximating model. With a fixed β , the increased persistence makes the agent save more than under the approximating model, which the observational equivalence proposition offsets by decreasing the household's patience via (9.3.18).

Fig. 9.5.1 and Fig. 9.5.2 record impulse response functions for the total endowment d_t under the approximating model and a worst case model associated with $\sigma = -.0001$, where β is held fixed at HST's benchmark value. Because these figures do not adjust the discount factor according to (9.3.18) as was done for Fig. 9.4.1 and Fig. 9.4.2, the distorted impulse response functions deviate from those of the approximating model even more than those of these earlier figures. The reduction in β from (9.3.18) works through two channels to make the $\sigma < 0$ decision rule equal to that for a $\sigma = 0$ rule: (1) it brings the distorted impulse response functions closer to those of the approximating model, and (2) more impatience combats the precautionary savings motive.

²⁶ XXX These figures are computed by hst4.m.

 $^{^{27}\,}$ The observational equivalence proposition makes the decision rules equivalent under the approximating model.



Figure 9.4.1: Response of total endowment d_t to innovation in 'permanent' component d_t^* under the approximating model (dotted line) and the distorted model associated with the worst case shock (dashed line) for the $\sigma = -.0001, \beta = \beta(\sigma)$ model.

9.5. Another view of precautionary savings

As an aid to interpret the precautionary savings motive inherent in our model, appendix B asserts another observational equivalence proposition. Theorem 9.B.1 takes a baseline case where $\beta R = 1$ and shows that in its effects on (c, i), activating a concern for robustness operates just like an *increase* in the discount factor. This result is useful because the $\beta R = 1$ case forms a benchmark in the permanent income literature (for example, see Hall (1978)). Theorem 9.B.1 shows that the effects of raising a concern for robustness by putting $\sigma < 0$ are replicated by simply raising β so that $\beta R > 1$.

To use this result to shed more light on how the precautionary motive manifests itself in the decision rule for consumption, we consider the important special case that $\delta = \lambda = \tilde{\delta} = 0$. Then $\mu_{st} = \mu_{ct} = b - c_t$ and the consumption Euler equation (9.3.2*d*) without a concern about robustness becomes

$$b - c_t = E_t [(\beta R) (b - c_{t+1})].$$



Figure 9.4.2: Response of total endowment d_t to innovation in 'transitory' component \hat{d}_t under the approximating model (solid line) and the distorted model associated with the worst case shock (dotted line) for the $\sigma = -.0001, \beta = \beta(\sigma)$ model.

If $\beta R > 1$, this equation implies that $b - c_t > E_t(b - c_{t+1})$ or

$$c_t < E_t c_{t+1},$$
 (9.5.1) ["nengtrick "]

so that the optimal policy is to make consumption grow on average.

Theorem 9.B.1 shows that when $\beta R = 1$, a concern about robustness ($\sigma < 0$) has the same effect on c_t, i_t as setting $\sigma = 0$ and setting a particular β for which $\beta R > 1$. Therefore, when $\beta R = 1$, the precautionary saving that occurs when $\sigma < 0$ is follows from (9.5.1). Activating a concern about robustness imparts an upward drift to the expected consumption profile.

We can also use Theorem 9.B.1 to say some things about the decision rule for consumption in our special case that $\lambda = \delta = \tilde{\delta} = 0$. The solution (9.3.8) for $\sigma = 0$ implies the consumption rule

$$c_t = \left(1 - R^{-2}\beta^{-1}\right) \left[Rk_{t-1} + E_t \sum_{j=0}^{\infty} R^{-j} d_{t+j} \right] + \left(\frac{(R\beta)^{-1} - 1}{R - 1}\right) b. \quad (9.5.2) \quad ["consfunction "]$$

Notice that the marginal propensity to consume out financial wealth Rk_{t-1} equals that out of human wealth $E_t \sum_{j=0}^{\infty} R^{-j} d_{t+j}$.²⁸ Further, an increase in β decreases the constant $+\left(\frac{(R\beta)^{-1}-1}{R-1}\right)b$ and increases the marginal propensity to consume $1 - R^{-2}\beta$. Relative to the baseline $\beta R = 1$ case, raising β raises the marginal propensity to consume out of wealth by $R^{-1}(1 - (R\beta)^{-1})$. This increase in the marginal propensity to consume still allows wealth to have an upward trajectory because of the reduction in the second term $\frac{(R\beta)^{-1}-1}{R-1}b$.

Tom: expand this transition

The following section views the precautionary savings motive from the frequency domain.



Figure 9.5.1: Response of total endowment d_t to innovation in 'permanent' component d_t^* under the approximating model (solid line) and the distorted model associated with the worst case shock (dotted line) for $\sigma = -.0001$, with β at benchmark value.

 $^{^{28}}$ This implication of precautionary savings coming from robustness differs from that coming from convex marginal utility functions, where precautionary savings reduces the marginal propensity to consume out of endowment income relative to that from financial wealth. See Wang (2002XXX).


Figure 9.5.2: Response of total endowment d_t to innovation in 'permanent' component d_t^* under the approximating model (solid line) and the distorted model associated with the worst case shock (dotted line) for $\sigma = -.0001$ with β at benchmark value.

9.6. Frequency domain representation

This section uses HST's estimated permanent income model to illustrate features of the frequency domain decompositions of the consumer's objective function and of the worst case shocks for different values of σ .

To import some notation from chapter 7, denote the transfer function from shocks ϵ_t to the 'target' $s_t - b_t$ as $G(\zeta)$. For the baseline model with habit persistence, recall formula (7.3.6) for the frequency decomposition of H_2

$$H_{2} = -\frac{1}{2\pi} \int_{-\pi}^{\pi} \operatorname{trace} \left[G\left(\sqrt{\beta} \exp\left(i\omega\right)\right)' G\left(\sqrt{\beta} \exp\left(i\omega\right)\right) \right] d\omega$$

A reinterpretation of formula (7.3.5) also gives us the frequency domain representation

$$E\sum_{t=0}^{\infty}\beta^{t}w_{t}'w_{t} = \frac{1}{2\pi}\int_{-\pi}^{\pi}W\left(\sqrt{\beta}\exp\left(i\omega\right)\right)'W\left(\sqrt{\beta}\exp\left(i\omega\right)\right)d\omega.$$

For the baseline ($\sigma = 0$) line,²⁹ Fig. 12.5.1 shows $G(\sqrt{\beta} \exp(i\omega))'G(\sqrt{\beta} \exp(i\omega))$ as a function of frequency ω ; G'G is larger at lower frequencies. Remember that $G(\zeta) = (I - (A_o - BF)\zeta)^{-1}C$ embodies the consumer's optimal decision rule F. The noise process ϵ_t upon which $G(\zeta)$ operates is i.i.d. under the approximating model, so that the spectral density matrix of ϵ_t is constant across frequencies. But seeing that the consumer's policy makes him most vulnerable to the low frequency components of ϵ_t , the minimizing player makes the conditional mean of the worst-case shock w_{t+1} highly serially correlated. For two values of σ , Fig. 12.5.2 shows frequency decompositions of trace $W(\zeta)'W(\zeta)$ for $\zeta = \sqrt{\beta} \exp(i\omega)$. Notice how most of the power is at the lowest frequencies. As we varied σ from zero to the two values in Fig. 12.5.2, we adjusted $\beta = \hat{\beta}$ according to (9.3.18), which keeps the robust $\sigma < 0$ decision rule for consumption equal that to that for the baseline no robustness ($\sigma = 0$) model. Notice that [trace $W(\zeta)'W(\zeta)$] varies directly with the absolute value of σ .



Figure 9.6.1: Frequency decomposition of criterion function; $G(\zeta)'G(\zeta)$ plotted as a function of ω where $\zeta = \sqrt{\beta} \exp(i\omega)$.

 $^{^{29}}$ XXXX These figures were computed by hst3.m in the hst directory.



Figure 9.6.2: Frequency decomposition of volatility of worst case shocks for $-\theta^{-1} = \sigma = -.0001$ (solid line) and $\sigma = -.00005$ (dotted line); trace[$W(\zeta)'W(\zeta)$] plotted as a function of ω where $\zeta = \sqrt{\beta} \exp(i\omega)$.



Figure 9.6.3: Detection error probabilities as a function of σ .

9.7. Detection error probabilities

For HST's parameter values, Fig. 9.6.3 reports detection error probabilities associated with various values of σ , adjusting β according to (9.3.18) to keep the decision rule fixed. These detection error probabilities were calculated by the method of chapter 8 for a sample of the same length that HST used to estimate their model and for HST's initial condition. To calculate the detection error probabilities, all other parameter values were frozen at the values in Table 9.A.1. Then the formula for the worst-case distortions $w_{t+1} = K(\sigma)x_t$ was used to compute an alternative law of motion for the endowment process.

For different values of σ , Fig. 9.6.3 records the detection error probabilities for distinguishing an approximating model from a worst-case model associated with that value of σ . The approximating model is

$$x_{t+1} = (A - BF(0)) x_t + C\epsilon_{t+1}$$

while the distorted model associated with σ is

$$x_{t+1} = \left(A - BF\left(0\right) + CK\left(\sigma\right)\right)x_t + C\tilde{\epsilon}_{t+1}$$

where both ϵ_t and $\tilde{\epsilon}_t$ are i.i.d. processes with mean zero and identity covariance matrix, and where $F(0) = F(\sigma)$ by the observational equivalence proposition.

The detection error probability equals .5 for $\sigma = 0$ because then the models are identical and so cannot be distinguished. The detection error probability falls with σ because the two models spread from one another. In the following section, we use Fig. 9.6.3 to guide a choice of σ as measuring the size of a set of models against which it is plausible for the consumer to want to be robust.

9.8. Robustness of decision rules

For $\sigma = -\theta^{-1}$, express the equilibrium decision rules of game (9.2.1) as

$$c_{t} = -F(\sigma) x_{t} \qquad (9.8.1a) \quad ["soln1;a"]$$

$$w_{t+1} = K(\sigma) x_t$$
 (9.8.1b) ["soln1;b"]

and express $s_t - b$ as $H(\sigma)x_t$. For possibly different values σ_1, σ_2 , consider the law of motion of the state under the consumption plan $F(\sigma_2)x_t$ and the worst case shock process $K(\sigma_1)x_t$:

$$x_{t+1} = (A - BF(\sigma_2) + CK(\sigma_1)) x_t + C\epsilon_{t+1}.$$
 (9.8.2) ["soln2 "]

For x_0 given, we evaluate the expected payoff

$$\pi\left(\sigma_{1};\sigma_{2}\right) = -E_{0,\sigma_{1}}\sum_{t=0}^{\infty}\beta^{t}x_{t}'H\left(\sigma_{2}\right)'H\left(\sigma_{2}\right)x_{t} \qquad (9.8.3) \quad \texttt{["soln3 "]}$$

under the law of motion (9.8.2). That is, we want to evaluate the performance of the rule designed by setting σ_2 when the data are generated by the distorted model associated with σ_1 . For three values of σ_2 , Fig. 9.6.3 plots $\pi(\sigma_1; \sigma_2)$ as a function of the parameter σ_1 that indexes the magnitude of the distortion in the model generating the data. By construction, the $\sigma_2 = 0$ does better than the other rules when $\sigma_1 = 0$. But its performance deteriorates faster with decreases in σ_1 below zero than do the more robust $\sigma_1 = -.00004, \sigma_1 = -.00008$ rules.

From Fig. 9.6.3, $\sigma = -.00004$ is associated with a detection error probability of over .3, and $\sigma = -.00008$ with a detection error probability about .2. It is plausible for the consumer to want decisions that are robust against alternative models as close as the worst case models associated with those values of σ .



Figure 9.8.1: $\pi(\sigma_1; \sigma_2) = -E_{0,\sigma_1} \sum_{t=0}^{\infty} \beta^t x'_t H(\sigma_2)' H(\sigma_2) x_t$ as a function of σ_1 on the ordinate axis for decision rules $F(\sigma_2)$ associated with three values of σ_2 .

9.9. Concluding remarks

Different observationally equivalent (σ, β) pairs identified by Theorem 9.3.1 bear different implications about (i) pricing risky assets; (ii) the amounts required to compensate the planner for confronting different amounts of risk; (iii) the amount of model misspecification used to justify the planner's decisions if risk sensitivity is reinterpreted as reflecting concerns about model misspecification. Hansen, Sargent, and Tallarini (1999) and Hansen, Sargent, and Wang (2002) have extracted some asset pricing implications of the model in this chapter. They show that although movements along the observational equivalence locus laid out by (9.3.18) don't affect consumption and investment, they put an adjustment for fear of model misspecification into asset prices and boost measured market prices of risk. In chapter 12, we shall describe how standard asset pricing formulas are altered when a representative agent is concerned about robustness. There we shall describe an asset pricing theory under a concern about robustness in the context of a class of general equilibrium models of which the model of this chapter can be viewed as a special case.

object	habit	no habit
	Persistence	Persistence
risk free rate	.025	.025
β	.997	.997
δ_h	.682	
λ	2.443	0
α_1	.813	.900
α_2	.189	.241
ϕ_1	.998	.995
ϕ_2	.704	.450
μ_d	13.710	13.594
$c_{\hat{d}}$.155	.173
c_{d^*}	.108	.098
$2 \times \text{LogLikel}$	779.05	762.55

Table 9.A.1: HST's parameter estimates

A. Parameter values

HST calibrated a $\sigma = 0$ version of their permanent income model by maximizing a likelihood function conditioned only on U.S. post-war quarterly consumption and investment data. They used U.S. quarterly data on consumption and investment for the period 1970I–1996III. They measured consumption by nondurables plus services and investment by the sum of durable consumption and gross private investment.³⁰ They estimated the model from data on (c_t, i_t) , setting $\sigma = 0$, then deduced pairs (σ, β) that are observationally equivalent, using formula (9.3.18).

The forcing processes are governed by seven free parameters: $(\alpha_1, \alpha_2, c_{\hat{d}}, \phi_1, \phi_2, c_{d^*}, \mu_d)$. The parameter μ_b set a bliss point. While μ_b alters the marginal utilities, it does not influence the decision rules for consumption and investment. HST fixed μ_b at an arbitrary number, namely 32, for estimation.

Four parameters govern the endogenous dynamics: $(\gamma, \delta_h, \beta, \lambda)$. HST set $\delta_k = .975$, and imposed the permanent income restriction, $\beta R = 1$. The restrictions that $\beta R = 1, \delta_k = .975$ pin down γ once β is estimated. HST imposed $\beta = .9971$, which after adjustment for the effects of the geometric growth factor of 1.0033 implies an annual real interest rate of 2.5%.

Table 9.A.1 reports HST's estimates for the parameters governing the endogenous and exogenous dynamics. Fig. 9.A.1 and Fig. 9.A.2 report impulse response functions for consumption and investment to innovations in both components of the endowment process. For comparison, Table 9.A.1 reports estimates from a no habit persistence $(\lambda = 0)$ model.

 $^{^{30}}$ They estimated the model from data that had been scaled through multiplication by 1.0033^{-t} .



Figure 9.A.2: Impulse response of investment (circles) and consumption (line) to innovation in persistent shock (d^*) , at maximum likelihood estimate of habit persistence.

Notice that the persistent endowment shock process contributes much more to consumption and investment fluctuations than does the transitory endowment shock process.



Figure 9.A.1: Impulse response of investment (circles) and consumption (line) to innovation in transitory endowment process (\hat{d}) , at maximum likelihood estimate of habit persistence.

B. Another observational equivalence result

To shed more light on the form of precautionary savings, we state another observational equivalence result that takes as its benchmark an initial allocation associated with parameter settings $\beta R = 1$ and $\sigma < 0$. Then we find another value of β that implies the same decisions for c_t, i_t as the base model when $\sigma = 0$, so that the decision maker fears model misspecification. This entails working backwards from the worst case model that is reflected in the $\sigma < 0$ decision rule to the associated approximating model.

Theorem 9.B.1. (Observational Equivalence, II) Fix all parameters except (β, σ) . Consider a consumption-investment allocation for $(\hat{\beta}, \hat{\sigma})$ where $\hat{\beta}$ satisfies $\hat{\beta}R = 1$ and $\hat{\sigma} < 0$ and $\hat{\sigma} < \hat{\sigma}$. Then there exists a $\tilde{\beta} > \hat{\beta}$ such that the $(\hat{\beta}, \hat{\sigma})$ allocation also solves the $(\tilde{\beta}, 0)$ problem.

Proof. We suppose that $\hat{\sigma} < 0$, so that the worst case model differs from the approximating model. We want to find the approximating model and a value $\tilde{\beta}$ of β for which a $\sigma = 0$ decision maker would choose the $\hat{\beta}, \hat{\sigma}$ allocation. Under the model with $\hat{\sigma} < 0$, where \hat{E}_t denotes a conditional expectation under the worst case model, we have

$$\hat{E}_t \mu_{c,t+1} = \mu_{c,t}$$
 (9.B.1) ["nobseq3 "]

because $\hat{\beta}R = 1$. Let

$$\hat{E}_t \mu_{s,t+1} = \xi \left(\tilde{\beta} \right) \mu_{s,t}. \tag{9.B.2} \quad ["nobseq4 "]$$

Equation (9.B.1) implies that we want

$$1 = \xi \left(\tilde{\beta} \right) \tag{9.B.3} ["nobseq5 "]$$

where the projection coefficient $\xi(\hat{\beta})$ emerges from the multiplier problem for the evil agent for $\hat{\sigma} < 0$, which can be cast as

$$\min_{\{w_{t+1}\}} \left[-\sum_{t=0}^{\infty} \hat{\beta}^t \{ \mu_{st}^2 + \hat{\beta} \frac{1}{\hat{\sigma}} w_{t+1}^2 \} \right]$$

subject to the law of motion

$$\mu_{st} = \delta\left(\tilde{\beta}\right)\mu_{s,t-1} + \alpha w_t \qquad (9.B.4) \quad ["nobseq6"]$$

where $\delta(\tilde{\beta}) = \frac{1}{\beta R}$ and α is given by (9.3.17), (9.3.14), (9.3.15) under the $(\hat{\beta}, \hat{\sigma})$ model. (Remember that the decision rule for c_t and therefore the law for μ_{st} will be the same under our two observationally equivalent β, σ pairs, so we can use the benchmark case to compute α .) We freeze all parameters except β, σ . The approximating model would be $\mu_{st} = \delta \mu_{s,t-1} + \alpha \epsilon_t$, so that (9.B.4) adds a perturbation αw_t to the law of motion of μ_{st} under a deterministic version of the approximating model. The Bellman equation for the minimizing agent is evidently

$$-P\mu_s^2 = -\mu_s^2 + \hat{\beta}\min_w \left[-\frac{1}{\hat{\sigma}} w^2 - P\left(\delta\mu_s + \alpha w\right)^2 \right].$$
(9.B.5) ["bellmanneng "]

Notice the presence of both $\hat{\beta}$ and $\tilde{\beta}$, via δ and α . The first-order condition is

$$w = K\mu_s,$$

where

$$K = -\frac{\alpha \delta \hat{\sigma} P}{1 + \alpha^2 \hat{\sigma} P}.$$

Notice that

$$\xi\left(\tilde{\beta}\right) = A + KC = \delta + K\alpha = 1$$

which implies that

$$1 = \xi\left(\tilde{\beta}\right) = \delta + K\alpha = \frac{\delta}{1 + \alpha^2 \hat{\sigma} P}.$$

 $\delta = 1 + \hat{\sigma}\alpha^2 P < 1.$

Therefore,

Equation (9.B.5) implies that

$$-P = -1 + \hat{\beta} \left[-\frac{1}{\hat{\sigma}} K^2 - P \left(\delta + K \alpha \right)^2 \right].$$

Simplifying the above identity leaves

$$P = \frac{1}{1 - \hat{\beta}} \left[1 + \frac{\hat{\beta}}{\hat{\sigma}} \left(\frac{1 - \delta}{\alpha} \right)^2 \right].$$
(9.B.7) ["P"]

Equations (9.B.6) and (9.B.7) together imply that

$$0 = \hat{\beta} \left(1 - \delta \left(\tilde{\beta} \right) \right)^2 + \left(1 - \hat{\beta} \right) \left(1 - \delta \left(\tilde{\beta} \right) \right) + \alpha \left(\tilde{\beta} \right)^2 \hat{\sigma}.$$

A solution of this equation determines $\tilde{\beta}$. The solution of this quadratic equation is

$$\delta = 1 - \frac{-\left(1 - \hat{\beta}\right) \pm \sqrt{\left(1 - \hat{\beta}\right)^2 - 4\hat{\beta}\sigma\alpha^2}}{2\hat{\beta}}.$$

If $\sigma = 0$, this equation implies $\delta = 1$. When $\sigma < 0$, the appropriate root is

$$\delta = 1 - \frac{-\left(1 - \hat{\beta}\right) + \sqrt{\left(1 - \hat{\beta}\right)^2 - 4\hat{\beta}\sigma\alpha^2}}{2\hat{\beta}}.$$

Using $\hat{\beta}R = 1$, this is equivalent to

$$\tilde{\beta}\left(\sigma\right) = \frac{\hat{\beta}\left(1+\hat{\beta}\right)}{2\left(1+\sigma\alpha^{2}\right)} \left[1+\sqrt{1-4\hat{\beta}\frac{1+\sigma\alpha^{2}}{\left(1+\hat{\beta}\right)^{2}}}\right].$$
(9.B.8) ["obsequivn2 "]

Chapter 10 Competitive equilibrium models

10.1. Introduction

The next three chapters study prices and quantities in a dynamic competitive equilibrium model when a representative agent fears model misspecification. This chapter sets the stage by describing competitive equilibria when the representative agent has no concern about model misspecification. It introduces the basic objects and equilibrium representations of prices and quantities that will be modified when we add concerns about model misspecification in chapters 11 and 12.

10.2. Pricing risky claims

In an economy with complete markets, history-date prices equal intertemporal marginal rates of substitution times conditional probabilities evaluated at an equilibrium allocation. Complete markets assure that intertemporal rates of substitution are equated across all consumers, making it possible to speak unambiguously of *the* intertemporal rate of substitution and thereby allowing us to synthesize a representative agent.¹

In a pure endowment economy that directly specifies the preferences of a representative consumer, like the economy studied by Robert E. Lucas, Jr. (1978), it is trivial to compute the equilibrium history-date prices. They equal the representative consumer's intertemporal marginal rates of substitution, evaluated at the endowment, times the exogenous conditional probabilities. These

– 251 –

¹ Hansen and Sargent (XXXX, chapter YYYY) extend methods for calculating a representative household in a heterogenous agent economy to situations where households can have different dynamic technologies for transforming consumption goods into household services (e.g., different degrees of habit persistence).

prices can be used to evaluate risky claims that consist of bundles of history-date contingent commodities.

Brock (1982XXX) extended this pricing strategy to representative household economies that have endogenous state variables such as capital stocks that produce goods and household services. Household capital stocks can be used to represent non-separabilities over time in the household's preferences. Brock's procedure for pricing risky claims has these steps:

- 1. Compute optimal allocations by solving a planning problem. The optimal allocations can be represented recursively as functions of a state vector, the endogenous components of which are influenced by the planner's decisions.
- 2. Compute shadow prices of the history-date contingent consumption goods as conditional probabilities times intertemporal marginal rates of substitution evaluated at the optimal allocation. Take the shadow prices to be the history-date contingent prices.
- 3. Represent a security as a stochastic process of pay outs, i.e., as a sequence of measurable functions of the economy's history of shocks.
- 4. Price a security by multiplying the history-date payouts by the history-date prices computed in step (2), then sum over time and across histories.

Chapter 12 relies heavily on the four step strategy (1)-(4) for pricing assets. Here and in chapter 11, we lay out alternative decentralizations of our planning problem. Our four step strategy is powerful partly because it allows us to price assets using the state vector x_t for the planning problem. However, to express the idea that the household is a price taker in a recursive competitive equilibrium, we have to augment the state x_t with additional components X_t that are comparable in dimension to x_t ; X_t becomes a part of the state vector that the household takes as exogenous and in terms of which we express prices.² In a competitive equilibrium, we impose $X_t = x_t$, but only after the household has optimized while taking X_t as beyond its control. Setting $X_t = x_t$ after optimization is what makes the representative household be representative. After we have set $X_t = x_t$, we can cast asset pricing formulas solely in terms of the state x_t in the planning problem, provided that we adopt the assumption of time-zero trading that is embedded in the standard Arrow-Debreu model of

 $^{^2\,}$ In related contexts, this idea was used by Lucas and Prescott (1971) and Prescott and Mehra (1982XXX).

competitive equilibrium. The Arrow-Debreu timing is contained in one of the three types of competitive equilibrium models to be described in this chapter.

10.3. Types of competitive equilibria

We study a class of economic environments that fit the optimal linear regulator.³ Three types of competitive equilibrium share common specifications of information, preferences, and technologies, but have different market structures. They are (1) an "Arrow-Debreu equilibrium" with trades at time 0 in a complete set of state-contingent dated commodities; (2) an "equilibrium with Arrow securities" that has a sequence of complete markets in current period commodities and one-period ahead state-contingent claims; and (3) a "partial equilibrium" model in which a competitive representative firm acts as a price taker and prices lie along a system of demand equations perturbed by shocks. The allocations in all of the competitive equilibria provide alternative decentralized ways of attaining the same allocation by confronting households and firms with different price systems and trading opportunities. For applications, it is useful to know how to transform one type of equilibrium into another.

10.4. Information, preferences, and technology

 $^{^{3}}$ Hansen (1987) and Hansen and Sargent (200X) have studied such economies.

10.4.1. Information

An exogenous information vector z_t is governed by

$$z_{t+1} = A_{22}z_t + C_2\epsilon_{t+1},\tag{10.4.1}$$

where $\{\epsilon_t\}$ is an i.i.d. Gaussian vector with mean 0 and covariance matrix I, and the eigenvalues of $\tilde{A}_{22} \equiv \sqrt{\beta}A_{22}$ are bounded by unity in modulus. The vector z_t determines a time t preference shock b_t and a time t endowment shock d_t via

$$d_t = U_d \ z_t$$

$$b_t = U_b \ z_t.$$
(10.4.2)

To account for the flow of information in the economy, we define the space $J_t = [\epsilon^t, x_0]$, where $J_0 = [x_0]$ and $\epsilon^t = [\epsilon_t, \epsilon_{t-1}, \ldots, \epsilon_1]$. We say that a stochastic process is 'adapted to J_t ' if its time t component is a measurable function of J_t .

10.4.2. Preferences

A representative household has preferences ordered by

$$-(1/2) E\left(\sum_{t=0}^{\infty} \beta^t \left(|s_t - b_t|^2 + \ell_t^2\right) \mid J_0\right), \qquad (10.4.3)$$

where ℓ_t is a scalar process that constrains a vector g_t of intermediate activities (designed to capture generalized adjustment costs) in equation (10.4.6) below, and s_t is a vector of household services produced at time t via the household technology

$$s_t = \Lambda h_{t-1} + \Pi c_t$$

$$h_t = \Delta_h h_{t-1} + \Theta_h c_t.$$
(10.4.4)

Sometimes we interpret ℓ_t in (10.4.3) as labor input. In (10.4.4), h_t is a vector of stocks of household durable goods at t, c_t is a vector of consumption flows, and Λ , Π , Δ_h , Θ_h are matrices.

10.4.3. Technology

There is a constant returns to scale production technology

$$\Phi_c c_t + \Phi_i i_t + \Phi_g g_t = \Gamma k_{t-1} + d_t$$

$$k_t = \Delta_k k_{t-1} + \Theta_k i_t,$$
(10.4.5)

where k_t is a vector of capital goods used in production, i_t is a vector of investment goods, Δ_k is a matrix, and g_t is constrained by⁴

$$g_t \cdot g_t \le \ell_t^2. \tag{10.4.6}$$

10.4.4. Planning problem

The planning problem is to maximize (10.4.3) over choices of processes for $\{s_t, c_t, i_t, g_t, k_t, h_t\}_{t=0}^{\infty}$ that are adapted to J_t subject to (10.4.1), (10.4.2), (10.4.4), and (10.4.5) with given initial conditions for (z_0, h_{-1}, k_{-1}) . The planning problem takes the form of an optimal linear regulator. Let

$$x_t \equiv \begin{bmatrix} h_{t-1} \\ k_{t-1} \\ z_t \end{bmatrix}.$$

The two components h and k of the state vector are *endogenous* and z is *exogenous*. If the matrix $\Phi \equiv [\Phi_c \quad \Phi_g]$ is nonsingular, the control vector u_t can be chosen to be investment i_t because⁵

$$\begin{bmatrix} c_t \\ g_t \end{bmatrix} = \Phi^{-1} \left(\Gamma k_{t-1} + U_d z_t - \Phi_i i_t \right).$$
 (10.4.7)

Using this relation, the constraints (10.4.4) and (10.4.5) can be rewritten

$$x_{t+1} = Ax_t + Bu_t + C\epsilon_{t+1} \tag{10.4.8}$$

⁴ Under the constant returns to scale interpretation, d_t is taken as an additional input available in fixed supply.

⁵ The matrix Φ can usually be rendered nonsingular by augmenting the control vector to include some components of consumption or the labor-using intermediate activities.

for appropriately chosen matrices A, B, C. The matrix A is block triangular and the bottom row block of B is zero as required for the discounted stochastic linear regulator problem (see page 64). Moreover, using (10.4.7) and (10.4.4), the time t terms $|s_t - b_t|^2$ and $|g_t|^2$ in the objective function (10.4.3) of the planner can both be expressed as quadratic forms in the control $u_t = i_t$ and the augmented state x_t .

The planner's optimal decision rule is $u_t = -Fx_t$. Under this rule, the state evolves according to

$$x_{t+1} = A^o x_t + C\epsilon_{t+1}, (10.4.9)$$

where $A^o = A - BF$.

10.4.5. Imposing stability

In permanent income economies, optimality does not automatically imply stability of the state vector process. For example, the economy of chapter 9 has a single consumption good, a single capital good, and no labor-using intermediate activities g_t . The counterpart to equation (10.4.7) is

$$c_t = \Gamma k_{t-1} + U_d z_t - i_t.$$

The chapter 9 model constrains the subjective discount factor to be the reciprocal of the physical return to capital: $\beta = \frac{1}{\Gamma + \Delta_k}$. Without imposing stability as an additional constraint, the optimal sequence of capital stocks diverges to minus infinity at a rate that is not dominated by $\frac{1}{\sqrt{\beta}}$. We want to impose stability because solutions that require x_t not to explode at a rate exceeding $\frac{1}{\sqrt{\beta}}$ are much better approximations to models that impose debt limits or various non-negativity constraints. We therefore impose stability as an additional constraint, with the consequence that the solution of the resulting infinite-horizon control problem equals the limit of a sequence of solutions to the corresponding finite-horizon problems, each of which imposes a zero terminal capital stock.

We now describe competitive equilibria with three different types of trading structures, each of which supports an allocation that solves the planning problem.

10.5. Arrow-Debreu equilibrium

10.5.1. The price system at time 0

An Arrow-Debreu equilibrium has complete markets at time 0 in claims to history-contingent dated commodities. We follow Harrison and Kreps (1979) and Hansen and Sargent (XXXX) in re-scaling the Arrow-Debreu prices. The re-scaled prices are the ordinary Arrow-Debreu history-date prices divided by probabilities times discount factors. Using the scaled prices converts present values into expected discounted geometric sums of quadratic forms that are easy to compute by solving Sylvester equations (see chapter 3).

We use a price system with components $\{p_{ct}^0, p_{it}^0, p_{\ell t}^0, p_{dt}^0, p_{rt}^0\}_{t=0}^{\infty}$, each element of which resides in a space L_0^2 defined by

$$L_0^2 = \left[\{y_t\}_{t=0}^\infty : y_t \text{ is a random variable in } J_t \text{ for } t \ge 0, \text{ and } E\left[\sum_{t=0}^\infty \beta^t y_t^2 \mid J_0\right] < +\infty \right]$$

That ' y_t is in J_t ' means that y_t can be expressed as a measurable function of J_t . The square summability requirement, $E[\sum_{t=0}^{\infty} \beta^t y_t^2 \mid J_0] < \infty$, imposes that y_t not grow too fast in absolute value.

Our price system contains the following prices: p_{ct}^0 is an $n_c \times 1$ stochastic process that prices the consumption process c_t ; $p_{\ell t}^0$ is a scalar stochastic process that prices ℓ_t ; p_{dt}^0 is a vector stochastic process that prices the process $\{d_t\}$; p_{it}^0 is an $n_k \times 1$ vector stochastic process that prices new investment goods; and p_{rt}^0 is an $n_k \times 1$ vector stochastic process of capital rental rates. A time tcomponent of the price system is a random vector that is a function of J_t . The *price system* is a *sequence* of vector-valued measurable functions of the time thistories J_t .

Prices and quantities are stochastic processes. We require the stochastic processes for both prices and quantities to reside in L_0^2 . By virtue of a Cauchy-Schwartz inequality, this makes the conditional inner products to be used in the budget constraints and objective functions below well defined and finite in equilibrium. Later it will be convenient to obtain recursive representations for both prices and quantities.

We now describe the choice problems faced by a household and a firm within a competitive equilibrium in which all trades occur at time 0. The household and firm act as price takers. The allocations chosen by the household and the firm must be "realizable" in the sense that time t decisions depend only on information available at time t, i.e., they must reside in L_0^2 .

10.5.2. The household

We let E denote the mathematical expectation evaluated with respect to the joint probability distribution of $[\epsilon^t, x_0]$. We also let E_t denote $E(\cdot|J_t)$. The household chooses stochastic processes for $\{c_t, s_t, h_t, \ell_t i_t, k_t\}_{t=0}^{\infty}$, each element of which is in L_0^2 , to maximize

$$-\frac{1}{2} E_0 \sum_{t=0}^{\infty} \beta^t \left[(s_t - b_t) \cdot (s_t - b_t) + \ell_t^2 \right]$$
(10.5.1)

subject to

$$E\sum_{t=0}^{\infty} \beta^{t} \left(p_{ct}^{0} \cdot c_{t} + p_{it}^{0} \cdot i_{t} \right) \mid J_{0}$$

= $E\sum_{t=0}^{\infty} \beta^{t} \left(p_{\ell t}^{0} \ell_{t} + p_{rt}^{0} \cdot k_{t-1} + p_{dt}^{0} \cdot d_{t} \right) \mid J_{0}$ (10.5.2*a*)

$$s_t = \Lambda h_{t-1} + \Pi c_t \tag{10.5.2b}$$

$$h_t = \Delta_h h_{t-1} + \Theta_h c_t \tag{10.5.2c}$$

$$k_t = \Delta_k k_{t-1} + \Theta_k i_t \tag{10.5.2d}$$

$$b_t = U_b z_t \tag{10.5.2e}$$

$$d_t = U_d z_t \tag{10.5.2f}$$

$$z_{t+1} = A_{22}z_t + C_2\epsilon_{t+1} \tag{10.5.2g}$$

with h_{-1}, k_{-1}, z_0 given.

10.5.3. The firm

A firm rents capital and labor and buys the realization of the endowment process d_t . It uses these inputs to produce consumption goods and investment goods that it sells to the household. The firm chooses stochastic processes for $\{c_t, i_t, k_t, \ell_t, g_t, d_t\}$, each element of which is in L_0^2 , to maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t \left(p_{ct}^0 \cdot c_t + p_{it}^0 \cdot i_t - p_{rt}^0 \cdot k_{t-1} - p_{\ell t}^0 \ell_t - p_{dt}^0 \cdot d_t \right)$$
(10.5.3)

subject to

$$\Phi_c c_t + \Phi_g g_t + \Phi_i i_t = \Gamma k_{t-1} + d_t \tag{10.5.4}$$

$$-\ell_t^2 + g_t \cdot g_t = 0. \tag{10.5.5}$$

10.5.4. Competitive equilibrium with time-zero trading

A competitive equilibrium has all trades being made at time 0.

Definition 10.5.1. A competitive equilibrium is a price system $\{p_{ct}^0, p_{dt}^0, p_{dt}^0, p_{dt}^0, p_{it}^0, p_{rt}^0\}_{t=0}^{\infty}$ and an allocation $\{c_t, i_t, s_t, k_t, h_t, \ell_t, g_t\}_{t=0}^{\infty}$ that satisfy the following conditions:

- a. The allocation and each component of the price system and reside in the space L_0^2 .
- b. Given the price system, the allocation solves the problems of the household and firm.

10.5.5. Equilibrium computation

A strategy for computing an equilibrium is first to solve the planning problem for equilibrium quantities, then to compute shadow prices that we transform into equilibrium prices.

The optimal linear regulator can be used to solve a planning problem. The optimal law of motion for the state x_t and the value function for the planning problem contain enough information to compute competitive equilibrium prices. Let V(x) = -x'Px - p be the optimal value of the planning problem starting from initial state $x = [h' \ k' \ z']'$. The Bellman equation for the planning problem is

$$-x'Px - p = \max_{c,i,g} \left\{ -.5 \left[(s-b) \cdot (s-b) + g \cdot g \right] + \beta E \left(-x^{*'}Px^{*} - p \right) \right\}$$
(10.5.6)

subject to the linear constraints

$$\Phi_{c}c + \Phi_{g}g + \Phi_{i}i = \Gamma k + d$$

$$k^{*} = \Delta_{k}k + \Theta_{k}i$$

$$h^{*} = \Delta_{h}h + \Theta_{h}c$$

$$s = \Lambda h + \Pi c$$

$$z^{*} = A_{22}z + C_{2}\epsilon$$

$$b = U_{b}z$$

$$d = U_{d}z,$$
(10.5.7)

where * denotes a next period value. The time-invariant character of the planning problem makes the optimal decision rules time invariant. Time t decision rules are linear in the state vector x_t . We denote these rules $c_t = S_c x_t, g_t = S_g x_t, h_t = S_h x_t, i_t = S_i x_t, k_t = S_k x_t, s_t = S_s x_t$.

The law of motion for the state vector is linear:

$$x_{t+1} = A^o x_t + C\epsilon_{t+1} \tag{10.5.8}$$

where

$$A^{o} \equiv \begin{bmatrix} A_{11}^{o} & A_{12}^{o} \\ 0 & A_{22} \end{bmatrix}, C \equiv \begin{bmatrix} 0 \\ C_{2} \end{bmatrix}.$$
 (10.5.9)

The partitioning of the A^o and C matrices is according to the endogenous state vector $\begin{bmatrix} h'_{t-1} & k'_{t-1} \end{bmatrix}'$ and the exogenous state vector z_t . The zero restriction

on the (2,1) partition of A^o reflects the fact that the exogenous component of the state vector at time t + 1 does not depend on the endogenous state vector at time t. The zero restriction on the first rows in the partition of C reflects the fact that the endogenous state vector at time t + 1 is predetermined (i.e., depends only on time t information). The contingency plans for h_t and k_t are embedded in the part of (10.5.9) that determines the endogenous state vector $[h'_t k'_t]'$ as a function of x_t . In particular,

$$\begin{bmatrix} S_h \\ S_k \end{bmatrix} = \begin{bmatrix} A_{11}^o & A_{12}^o \end{bmatrix}.$$
(10.5.10)

The planner's decision rules are recursive in the sense that the time t decision depends on the state vector at time t, which in turn depends on the state vector at time t - 1.

10.5.6. Shadow prices

Equilibrium prices can be found by appropriately reinterpreting shadow prices as prices. Formulas for shadow prices corresponding to the elements of the price system $\{p_{ct}^0, p_{dt}^0, p_{dt}^0, p_{it}^0, p_{t}^0\}_{t=0}^{\infty}$ can be extracted from A^o and the matrix Pin the quadratic form in the value function. Evaluating these shadow prices at the equilibrium allocation recovers prices.

The time t component of these shadow prices are linear functions of x_t . In particular, the vector of shadow prices P_{ct}^0 for c_t is given by⁶

$$p_{ct}^0 = M_c x_t$$

where

$$M_c = \Theta'_h M_h + \Pi' M_s.$$
 (10.5.11)

Here $M_h x_t$ is the shadow price of consumer durables and $M_s x_t$ is the shadow price of household services. These shadow prices satisfy

$$M_h x_t = E\left[\sum_{\tau=1}^{\infty} \beta^{\tau} \left(\Delta'_h\right)^{\tau-1} \Lambda' M_s x_{t+\tau} | J_t\right]$$
(10.5.12)

$$M_s x_t = (s_t - b_t), \qquad (10.5.13)$$

⁶ Quantities $M_j x_t$ emerge from derivatives of the planner's value function and are measured in units of marginal utility.

where the mathematical expectation is evaluated with respect to the model (10.5.8) and where $b_t = S_b x_t$ and $s_t = S_s x_t$ is the planner's solution for s_t as a function of x_t . Also, let the planner's solution for intermediate inputs g_t be $g_t = S_g x_t$. Hansen and Sargent (XXXX) show that these other components of the shadow price system have representations⁷

$$p_{ct}^{0} = M_{c}x_{t}$$

$$p_{it}^{0} = M_{i}x_{t}$$

$$p_{rt}^{0} = \Gamma' M_{d}x_{t}$$

$$p_{dt}^{0} = M_{d}x_{t}$$

$$p_{\ell t}^{0} = |M_{\ell}x_{t}|.$$
(10.5.14)

where Hansen and Sargent (200XXX) give the following formulas for the M_j 's in terms of P and A^o :

$$M_{k} = 2\beta [0 \ I \ 0] PA^{o}$$

$$M_{h} = 2\beta [I \ 0 \ 0] PA^{o}$$

$$M_{s} = (S_{b} - S_{s})$$

$$M_{d} = \begin{bmatrix} \Phi_{c}' \\ \Phi_{g}' \end{bmatrix}^{-1} \begin{bmatrix} \Theta_{h}' M_{h} + \Pi' M_{s} \\ -S_{g} \end{bmatrix}$$

$$M_{c} = \Theta_{h}' M_{h} + \Pi' M_{s}$$

$$M_{i} = \Theta_{k}' M_{k}.$$

$$M_{\ell} = S_{g}$$
(10.5.15)

Here the partitions $[0 \ I \ 0]$ and $[I \ 0 \ 0]$ are conformable with the partition $[h'_{t-1}, k'_{t-1}, z_t]'$ of x_t .

⁷ Hansen and Sargent also compute a shadow price of capital, which they show is $p_{kt}^0 = (\Gamma' M_d + \Delta'_k M_k) x_t$.

10.5.7. Recursive representation of time 0 prices

Formulas (10.5.8) and (10.5.14) imply that we can regard the price system as consisting of sequences of measurable functions of the histories $J_t = [\epsilon^t, x_0]$. Equations (10.5.8) and (10.5.14) give a recursive representation of this price system in terms of the planner's state vector x_t . Although this representation of the price system turns out to be very convenient for asset pricing, it is not an appropriate representation for posing a recursive version of the household's optimization problem in a competitive equilibrium. We can obtain a recursive representation of the price system that will serve this purpose by introducing an additional state vector designed to keep track of the histories J_t for $t \ge 0$. In particular, define a new state vector X_t with components H_{t-1}, K_{t-1} that have the same dimensions as h_{t-1}, k_{t-1} , respectively:

$$X_t = \begin{bmatrix} H_{t-1} \\ K_{t-1} \\ z_t \end{bmatrix}.$$

The exogenous state vector z_t is a common component of x_t and of X_t . Impose an initial condition

$$X_0 = \begin{bmatrix} H_{-1} \\ K_{-1} \\ z_0 \end{bmatrix} = \begin{bmatrix} h_{-1} \\ k_{-1} \\ z_0 \end{bmatrix} = x_0.$$

Take the law of motion for X to be

$$X_{t+1} = A^o X_t + C\epsilon_{t+1} \tag{10.5.16a}$$

where A^o is the same matrix that appears in the representation (10.4.9) for the evolution of x_t under the planner's optimal control. Then we can represent the shadow price system as

$$p_t^0 = MX_t. (10.5.16b)$$

What is the purpose of this "big X" representation for prices?⁸ First, note that by setting $X_0 = x_0$, we assure that (10.5.16) reproduces the planner's shadow prices. But by expressing them in terms of X_t rather than x_t , we

⁸ In their concept of a recursive competitive equilibrium, Prescott and Mehra (1980) distinguished between the market wide level of capital \underline{k} and the level chosen by an individual k in order to represent price-taking behavior. Our purpose is somewhat different than theirs, which was not to get a recursive representation of time-0 Arrow-Debreu prices.

make these prices depend only on a state variable that is beyond the control of the household. The role of X_t is to account for the history $J_t = [\epsilon^t, x_0]$. In a competitive equilibrium, we want households and firms to influence the evolution of h_t and k_t , the endogenous components of x_t , but still to be price takers. Therefore, in expressing the choices facing households and firms in a competitive equilibrium, we use X_t to provide a recursive representation of prices.

10.5.8. Recursive representation of household's problem

The household chooses an allocation to maximize (10.5.1) subject to (10.5.2) and the price system (10.5.16a), (10.5.16b), which the household regards as exogenous. Thus the household maximizes

$$E_0 \sum_{t=0}^{\infty} \beta^t \left\{ -.5 \left[(s_t - b_t) \cdot (s_t - b_t) + \ell_t^2 \right] \right\}$$
(10.5.17)

subject to

$$E\sum_{t=0}^{\infty} \beta^{t} \left(p_{ct}^{0} \cdot c_{t} + p_{it}^{0} \cdot i_{t} \right) \mid J_{0}$$

= $E\sum_{t=0}^{\infty} \beta^{t} \left(p_{\ell t}^{0} \ell_{t} + p_{rt}^{0} \cdot k_{t-1} + p_{dt}^{0} \cdot d_{t} \right) \mid J_{0}$ (10.5.18*a*)

$$s_{t} = \Lambda h_{t-1} + \Pi c_{t}$$
(10.5.18b)
$$h_{t} = \Delta_{b} h_{t-1} + \Theta_{b} c_{t}$$
(10.5.18c)

$$h_t = \Delta_h h_{t-1} + \Theta_h c_t \qquad (10.5.18c)$$

$$k_t = \Delta_h k_{t-1} + \Theta_h i_t \qquad (10.5.18d)$$

$$b_t = U_b z_t \tag{10.5.18e}$$

$$d_t = U_d z_t \tag{10.5.18} f)$$

$$p_t^0 = M_0 X_t \tag{10.5.18h}$$

$$X_{t+1} = A^0 X_t + C\epsilon_{t+1}.$$
 (10.5.18*i*)

with h_{-1}, k_{-1}, z_0 given. For a given Lagrange multiplier μ_0^w attached to (10.5.18*a*), problem (10.5.17), (10.5.18) takes the form of an optimal linear regulator. (See section 11.4 for details and also for a way to compute μ_0^w from the solution of the planning problem.)

The maximizing choice of the household makes the time t component of $a_t = \begin{bmatrix} c_t & i_t & ell_t \end{bmatrix}'$ a function of the composite state $\begin{bmatrix} h_{t-1} \\ k_{t-1} \\ x_t \end{bmatrix}$:

$$a_{t} = -\begin{bmatrix} f_{1} & f_{2} & f_{3} \end{bmatrix} \begin{bmatrix} h_{t-1} \\ k_{t-1} \\ X_{t} \end{bmatrix}.$$
 (10.5.19)

The solution of the planning problem makes a_t a function of the state X_t :

$$a_t = -FX_t.$$
 (10.5.20)

In an equilibrium, $x_t = X_t$ in (10.5.19), so that the following equality prevails for all $x_t = X_t$: s $- \begin{bmatrix} f_1 & f_2 & f_3 \end{bmatrix} \begin{bmatrix} h_{t-1} \\ k_{t-1} \\ X_t \end{bmatrix} = -FX_t$.

10.5.9. Units of prices and reopening markets

Prices have the units of time 0 marginal utilities of the representative agent. We can choose a numeraire to express prices in terms of one of the consumption goods. In particular, denote the time t marginal utility of the first consumption good $e_1u_{c,t}$ and assume that $e_1u_{c,t} \neq 0$ with probability one for all t. This assumption makes the first consumption good at time t a legitimate numeraire. We choose to express the price system at time 0 in units of the first consumption good. Therefore, we set

$$p_t^0 = \frac{Mx_t}{e_1 u_{c,0}}.$$
 (10.5.21)

More generally, for $t \ge \tau$ where $\tau \ge 0$, we could define a time τ price system⁹

$$p_t^{\tau} = \frac{M}{e_1 u_{c,\tau}} \tag{10.5.22}$$

To convert the tail of the time 0 price system for $t \ge 1$ to the time 1 price system we can use

$$p_{ct}^1 = p_{ct}^0 \frac{e_1 u_{c,0}}{e_1 u_{c,1}},$$

⁹ This time τ price system would prevail if we were to reopen markets at time τ , subject to appropriate initial conditions being inherited from earlier trading at time 0 prices.

and so on.

10.6. Sequential markets with Arrow securities

We noted in section 10.5.7 that the time 0 prices for an Arrow-Debreu time 0 competitive equilibrium have a recursive representation based on (10.5.16a), (10.5.16b). This fact makes it easy to construct equilibrium prices for trading in a sequence of one-period markets. Following Arrow (1954XXX), we can use (10.5.16) to form competitive equilibrium prices with sequential trading of all current dated commodities and one-period state-contingent claims to a composite commodity called wealth. In this setting, the decision problem of the household is recursive and that of the firm is static. Though the trading arrangement differs, the equilibrium allocation is the same as the one attained in the equilibrium of the Arrow-Debreu model with time 0 trading of all history-contingent commodities for all dates.

10.6.1. Arrow securities

To explain why it is natural to move from an equilibrium with time 0 trading to an equilibrium with sequential trading, represent the consumer's budget constraint (10.5.2) as

$$E_0 \sum_{t=1}^{\infty} \beta^t \left(p_{ct}^0 \cdot c_t + p_{it}^0 \cdot i_t - p_{\ell t}^0 \ell_t - p_{rt}^0 \cdot k_{t-1} - p_{dt}^0 \cdot d_t \right) + p_{c0}^0 \cdot c_0 + p_{i0}^0 \cdot i_0$$

= $p_{\ell 0}^0 \ell_0 + p_{r0}^0 \cdot k_{-1} + p_{d0}^0 \cdot d_0.$

Express the expected discounted sum on the left side as

$$E_{0}\beta\left(\frac{e_{1}u_{c,1}}{e_{1}u_{c,0}}\right)\sum_{t=1}^{\infty}\beta^{t-1}\left(p_{ct}^{1}\cdot c_{t}+p_{it}^{1}\cdot i_{t}-p_{\ell t}^{1}\ell_{t}-p_{rt}^{1}\cdot k_{t-1}-p_{dt}^{1}\cdot d_{t}\right)$$

$$\equiv E_{0}\beta\left(\frac{e_{1}u_{c,1}}{e_{1}u_{c,0}}\right)a_{1}\left(X_{1}\right)=\int q\left(X_{1}|X_{0}\right)a_{1}\left(X_{1}\right)dX_{1},$$

(10.6.1)

where $a_1 = a_1(X_1)$ measures wealth at time 1 in state X_1 in units of the time 1 consumption good, and the one-step ahead pricing kernel $q(X_1|X_0) \equiv$

 $\beta \frac{e_1 u_{c,1}}{e_1 u_{c,0}} f(X_1|X_0)$. Here $f(X_1|X_0)$ is the transition density of X defined by (10.5.16a). Use (10.6.1) to express the budget constraint as

$$\int q(X_1|X_0)a_1(X_1) dX_1 + p_{c0}^0 \cdot c_0 + p_{i0}^0 \cdot i_0$$

$$= p_{\ell 0}^0 \ell_0 + p_{d0}^0 \cdot d_0 + p_{r0}^0 \cdot k_{-1} + a_0(X_0)$$
(10.6.2)

where $a_0(X_0)$ is the value of the household's initial wealth, namely, the capital stock k_{-1} . More generally, take prices without superscripts to be denominated in units of time t consumption of the first good and write

$$\int q (X_{t+1}|X_t) a_{t+1} (X_{t+1}) dX_{t+1} + p_{ct} \cdot c_t + p_{it} \cdot i_t$$

$$= p_{\ell t} \ell_t + p_{dt} \cdot d_t + p_{rt} \cdot k_{t-1} + a_t (X_t)$$
(10.6.3)

where

$$q(X_{t+1}|X_t) = \beta \frac{e_1 u_{c,t+1}}{e_1 u_{c,t}} f(X_{t+1}|X_t)$$
(10.6.4)

is the kernel for pricing claims of the first consumption good at time t + 1 in terms of time t consumption of the first good.

The spot prices p_{t} are given by the appropriate time t components of our original time 0 prices defined in (10.5.16). Together with the pricing kernel defined as (10.6.4), they allow us to support the solution of the planning problem by a competitive equilibrium with sequential markets. Within that equilibrium, the problem of the household is recursive and the problem of the firm is static, as we now proceed to show.

10.6.1.1. The household's problem in the sequential equilibrium

The household's Bellman equation is

$$W(a_{t}, h_{t-1}, X_{t}) = \max_{c_{t}, i_{t}, \ell_{t}, a(X_{t+1})} \left\{ -\left(|s_{t} - b_{t}|^{2} + \ell_{t}^{2}\right) + \beta \int W(a(X_{t+1}), h_{t}, k_{t}, X_{t+1}) f(X_{t+1}|X_{t}) dX_{t+1} \right\}$$
(10.6.5)

where the maximization is subject to

$$s_t = \Lambda h_t + \Pi c_t \tag{10.6.6a}$$

$$h_t = \Delta_h h_{t-1} + \Theta_h c_t \tag{10.6.6b}$$

$$k_t = \Delta_k k_{t-1} + \Theta_k i_t \tag{10.6.6c}$$

$$X_{t+1} = A^o X_t + C\epsilon_{t+1} \tag{10.6.6d}$$

$$p_t = MX_t \tag{10.6.6e}$$

$$b_t = S_b X_t \tag{10.6.6f}$$

$$d_t = S_d X_t \tag{10.6.6g}$$

$$\int a(X_{t+1}) q(X_{t+1}|X_t) dX_{t+1} = a_t + p_{\ell t} \ell_t + p_{dt} \cdot d_t + p_{rt} \cdot k_{t-1} - p_{ct} \cdot c_t - p_{it} \cdot i_t, \qquad (10.6.6h)$$

and $b_t = S_b X_t \equiv U_b z_t$. The optimal policy functions express c_t, i_t, ℓ_t , and $a(X_{t+1})$ each as functions of $(a_t, h_{t-1}, k_{t-1}, X_t)$.

10.6.1.2. The firm

The problem of a firm is static:

$$\max_{c_t, i_t, \ell_t, g_t, d_t, k_{t-1}} \left(p_{ct} \cdot c_t + p_{it} \cdot i_t - p_{rt} \cdot k_{t-1} - p_{dt} \cdot d_t - p_{\ell t} \ell_t \right)$$
(10.6.7)

subject to the technology (10.5.4), (10.5.5).

10.6.2. Recursive competitive equilibrium

We call a competitive equilibrium with Arrow securities a recursive competitive equilibrium. In a recursive competitive equilibrium, the household takes the law of motion for X_t as given. However, the household and the firm choose elements h_s, k_s of the state x_{s+1} that correspond to the elements H_s, K_s of X_{s+1} . A recursive competitive equilibrium requires that $X_t = x_t$ for $t \ge 1$, starting from $x_0 = X_0$, which means that the laws of motion chosen by firms and the household must be consistent with the law of motion (10.5.16) that inspires the household's decisions.

It can be verified that the quantities that solve the planning problem are recursive competitive equilibrium quantities at the candidate prices described above. See Hansen and Sargent (20XXX) for some of the details.

10.7. Asset pricing in a nutshell

It is a significant practical convenience that we can dispense with X_t as a state variable when we actually compute asset prices. After we have computed the equilibrium quantities and prices, for the purpose of computing asset prices, it is sufficient to express streams of payouts and prices both as functions of the state vector x_t that appears in the planning problem.¹⁰ As an example, let an endowment shock be a linear function of the exogenous component z_t , and let the endogenous component h_{t-1} track movements in the intertemporal rate of substitution that drives the prices. Let $\{y_t\}_{t=0}^{\infty}$ be a stochastic process of 'dividends', i.e., claims on the vector of consumption goods with representation $y_t = S_c x_t$. In units of the first time t consumption good, let a_{yt} denote the price at time t of a claim on the tail of the dividend process $\{y_s\}_{s=t}^{\infty}$. The price a_{yt} of a claim on the dividend stream from t onward can be represented as

$$x_{t+1} = A^o x_t + C\epsilon_{t+1} \tag{10.7.1a}$$

$$y_t = S_y x_t \tag{10.7.1b}$$

$$p_{cs} = (e_1 M_c x_t)^{-1} M_c x_s \tag{10.7.1c}$$

$$a_{yt} = E_t \sum_{t=0}^{\infty} \beta^t p_{ct+j} \cdot y_{t+j}.$$
 (10.7.1d)

Equation (10.7.1d) can be evaluated by solving a Sylvester equation.

Equations (10.7.1) capture the spirit of Brock's (1982) extension of Lucas's asset pricing formulas in which equilibrium history-date prices are the pertinent probabilities times the intertemporal marginal rate of substitution evaluated at the equilibrium allocation. The single state vector x_t tracks both the history-date prices and the dividend process.

¹⁰ We view this as a manifestation of Brock's (1982) idea of evaluating history-date prices by multiplying the relevant conditional probabilities by the representative consumer's marginal rates of substitution evaluated at the allocation that solves the planning problem. Incidentally, another perspective on Brock's insight is Mehra and Prescott's (1985) observation that the cross-equation restrictions between asset prices and consumption are not affected by whether consumption is regarded as exogenous or endogenous. Making consumption endogenous adds restrictions across consumption and yet other processes.

10.8. Partial equilibrium interpretation

Another decentralization of the planning problem makes contact with partial equilibrium models in the style of Lucas and Prescott (1971), Rosen, Murphy, and Scheinkman (1994), Rosen and Topel (199??), Ryoo and Rosen (2003), and Sargent (1987, chapter XIV). These models have a representative firm that acts as a price taker within an industry that faces a stochastically shifting linear demand schedule.

Within the environment of this chapter, consider a representative firm that chooses stochastic processes $\{c_t, g_t\}$ to maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t \left\{ p_t \cdot c_t - g_t \cdot g_t \right\}$$
(10.8.1)

subject to the constraints

$$k_t = \Delta_k k_{t-1} + \Theta_k i_t \tag{10.8.2a}$$

$$\Phi_c c_t + \Phi_g g_t + \Phi_i i_t = \Gamma k_{t-1} + d_t \tag{10.8.2b}$$

$$X_{t+1} = A^o X_t + C\epsilon_{t+1} \tag{10.8.2c}$$

$$d_t = U_d X_t \tag{10.8.2d}$$

$$p_t = M_c X_t. \tag{10.8.2e}$$

Here (10.8.2c), (10.8.2e) are used to represent a dynamic demand curve, where M_c and the state X_t are defined as above. The H_t component of X_t can express high order dynamics in demand.¹¹ The Bellman equation is

$$V(k_{t-1}, X_t) = \max_{c_t, k_t, i_t, g_t} \{ p_t \cdot c_t - g_t \cdot g_t + \beta EV(k_t, X_{t+1}) \}$$
(10.8.3)

where the maximization is subject to (10.8.2). The optimal decision rule expresses c_t, g_t as functions of k_t, X_t , so that the firm chooses to make k_t follow a law that can be expressed as

$$k_t = k(k_{t-1}, X_t) = k(k_{t-1}, H_{t-1}, K_{t-1}, z_t).$$
(10.8.4)

Embedded in (10.8.2c) is the firm's perceived law of motion for K_t , namely,

$$K_{t} = K(X_{t}) = K(H_{t-1}, K_{t-1}, z_{t}).$$
(10.8.5)

¹¹ See Hansen and Sargent (XXXX, chapters XXX) for an analysis.

A competitive equilibrium requires that $k_s \equiv K_s$ for all s and all z_s , or

$$k(K_{t-1}, H_{t-1}, K_{t-1}, z_t) = K(H_{t-1}, K_{t-1}, z_t).$$
(10.8.6)

The left side of (10.8.6) is the actual law of motion for K that emerges from *optimization* (this is the content of the function $k(\cdot)$) and *equilibrium* (this is the content of the condition k = K that makes the representative firm representative). The right side of (10.8.6) is the representative firm's perceived law of motion for K. Thus, (10.8.6) imposes equality between the 'perceived' law of motion for K and the 'actual' law of motion implied by those perceptions. Equality between these two laws imposes rational expectations and respects the price taking behavior of the firm.¹²

10.9. Concluding remarks

This chapter has set forth a class of dynamic linear quadratic economies and described three types of competitive equilibria. We have made the standard rational expectations assumption that a planner and the agents all trust their common model. The next chapter alters that assumption by instilling in both the household and a planner the same degree of preference for robustness of decision rules with respect to deviations of the actual data generation mechanism from a common approximating model. Thus, the next chapter will contain modifications of the Bellman equation (10.5.6), (10.5.7) for the planner and the Bellman equation (10.6.5), (10.6.6) for the household trading Arrow securities. With a preference for robustness, these will be replaced by Bellman equations for two-player, zero-sum games. Decentralizing the robust planning problem will then require checking that the choices of both the maximizing and the minimizing players for the planning problem and the household within a competitive equilibrium, respectively, are mutually consistent. Here agents choose an allocation as well as worst case shocks. To decentralize the economy, we shall require that these worst case shocks be appropriately aligned with those chosen by the planner.

¹² Marcet and Sargent (1989) and Evans and Honkapohja (2001) extensively exploit the definition of a rational expectations equilibrium as the fixed point of a mapping from a perceived to an actual law of motion.

Chapter 11 Competitive equilibrium under robustness

11.1. Introduction

This chapter puts fear of model misspecification into the mind of the representative household in the model of chapter 10. The analysis parallels that of chapter 10 except that we replace Bellman equations for the maximizing agents in chapter 10 with ones that pertain to two-player zero-sum games. In each game, a minimizing player helps the decision maker explore the fragility of a decision rule with respect to various difficult-to-detect perturbations of the approximating model. To decentralize an allocation that solves the planning problem, we specify two-player zero sum games for the representative household in a decentralized economy and describe how at the equilibrium prices both the allocation and the worst case model chosen by the household are aligned with counterparts that are chosen by the planner. Equilibrium quantities solve a robust planning problem. Competitive equilibrium prices can be computed from shadow prices for the robust planning problem. Chapter 12 uses these results to price assets when agents fear model misspecification. Appendix A describes a partial equilibrium decentralization of a robust planning problem with adjustment costs.

Before describing competitive equilibria with the more general household and production technologies of chapter 10, we begin with a pure endowment economy.

11.2. A pure endowment economy

11.2.1. The planning problem

We consider a nonstochastic pure endowment economy. Two sequences b_t, d_t are governed by

$$x_{t+1} = Ax_t + Cw_{t+1}$$
 (11.2.1*a*) ["simple1;a"]

$$b_t = U_b x_t \qquad (11.2.1b) \quad ["simple1;b"]$$

$$d_t = U_d x_t.$$
 (11.2.1c) ["simple1;c "]

Here x_t, b_t, d_t, w_t are each vectors; d_t is an exogenous endowment and b_t is an exogenous preference shock sequence; w_{t+1} is a distortion to the state x_t that determines the endowment and preference shock processes. A representative household receives consumption $c_t = d_t$. Because the endowment is exogenous, the robust planning problem becomes

$$-x_0' P x_0 = \min_{\{w_{t+1}\}} \sum_{t=0}^{\infty} \beta^t \left\{ -|c_t - b_t|^2 + \beta \theta w_{t+1}' w_{t+1} \right\}$$
(11.2.2) ["simple2 "]

where $\theta \in (\underline{\theta}, +\infty]$ measures the planner's fear that (11.2.1a) is misspecified. The planner minimizes criterion (11.2.2) subject to (11.2.1) and $c_t = d_t$. The planning problem is an optimal linear regulator. Let $R = (U_d - U_b)'(U_d - U_b)$. The Bellman equation for the planning problem is

$$-x'Px = \min_{w} \{ -x'Rx + w'(\beta\theta I)w - y'Py \}$$
(11.2.3*a*) ["simple3;a"]

where the minimization is subject to

$$y = Ax + Cw.$$
 (11.2.3b) ["simple3;b"]

The solution of the planning problem is a feedback rule

$$w = Kx.$$
 (11.2.4) ["simple4 "]

The law of motion for the state x under the worst case rule is

$$x_{t+1} = (A + KC) x_t$$

and the worst-case endowment and preference sequences can be represented

$$c_t = U_d (A + CK)^t x_0$$
 (11.2.5*a*) ["simple5;a"]

$$b_t = U_b (A + CK)^t x_0.$$
 (11.2.5b) ["simple5;b"]

The worst case processes in (11.2.5) generate a shadow price sequence $q_t = \beta^t p_t$ associated with the robust planning problem, where

$$p_t = M x_t$$
 (11.2.6) ["simple6 "]

and where $M = 2(U_b - U_d)$ or

$$p_t = 2(U_b - U_d)(A + CK)^t x_0.$$
 (11.2.7) ["simple7 "]

11.2.2. Household problem

We confront the representative household with the sequence of scaled (Harrison-Kreps) prices $\{p_t\}_{t=0}^{\infty}$ from (11.2.7). It is convenient to represent the price system recursively as

$$X_{t+1} = (A + CK) X_t$$
 (11.2.8*a*) ["simplep;a "]

$$p_t = MX_t \qquad (11.2.8b) \quad ["simplep;b"]$$

where $X_0 = x_0$. The household's problem is

$$\min_{\{w_{t+1}\}_{t=0}^{\infty}} \max_{\{c_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \{-|c_t - b_t|^2 + \beta \theta w_{t+1}' w_{t+1}\}$$
(11.2.9) ["simple8 "]

subject to (11.2.1) and the budget constraint

$$\sum_{t=0}^{\infty} \beta^t \left[p_t \cdot d_t - p_t \cdot c_t \right] = 0$$
 (11.2.10) ["simple9 "]

To solve this problem, we formulate a Lagrangian

$$\begin{split} L &= \sum_{t=0}^{\infty} \beta^t \{ -|c_t - b_t|^2 + \beta \theta w'_{t+1} w_{t+1} \\ &+ \mu_0^w [p_t \cdot d_t - p_t \cdot c_t] \} \end{split} \tag{11.2.11} \quad \texttt{["simple10 "]}$$

where μ_0^w is a Lagrange multiplier on (11.2.10). Formula (11.4.11) below implies that $\mu_0^w = 1$. We are to maximize (11.2.11) with respect to c_t and minimize it with respect to w_{t+1} . The first-order condition with respect to c_t is static:

$$\mu_0^w p_t = 2(b_t - c_t)$$
 (11.2.12) ["simple11 "]

which using $\mu_0^2 = 1$ leads to the demand curve

$$c_t = b_t - .5p_t.$$
 (11.2.13) ["simple12 "]

Using (11.2.12) or (11.2.13) and (11.2.8b), we deduce that $-|c_t-b_t|^2 = -.25X'_tM'MX_t$, $p_t \cdot d_t = X'_tM'U_dx_t$, and $p_t \cdot c_t = X'_tM'U_bx_t - .5X'_tM'MX_t$. Therefore the current period term in the objective (11.2.11) becomes $.25X_tM'MX_t + w'_{t+1}(\beta\theta I)w_{t+1} + X'_tM'(U_d - U_b)x_t$. Note that $M'(U_d - U_b) = -.5M'M$. Thus, we have that the Bellman equation for the household's problem reduces to

$$\begin{bmatrix} X\\x \end{bmatrix}' \begin{bmatrix} P_{11} & P_{12}\\P_{21} & P_{22} \end{bmatrix} \begin{bmatrix} X\\x \end{bmatrix} = \min_{w} \left\{ \begin{bmatrix} X\\x \end{bmatrix}' \begin{bmatrix} .25M'M & -.25M'M\\-.25M'M & 0 \end{bmatrix} \begin{bmatrix} X\\x \end{bmatrix} + w'(\beta\theta I)w + \beta \begin{bmatrix} Y\\y \end{bmatrix}' \begin{bmatrix} P_{11} & P_{12}\\P_{21} & P_{22} \end{bmatrix} \begin{bmatrix} Y\\y \end{bmatrix} \right\}$$
(11.2.14) ["simple13"]

where the minimization is subject to

$$\begin{bmatrix} Y \\ y \end{bmatrix} = \begin{bmatrix} A + CK & 0 \\ 0 & A \end{bmatrix} \begin{bmatrix} X \\ x \end{bmatrix} + \begin{bmatrix} 0 \\ C \end{bmatrix} w.$$
(11.2.15) ["simple14"]

The solution of the robust household's Bellman equation (11.2.14), (11.2.15) is attained by a policy rule

$$w = \begin{bmatrix} \tilde{K}_1 & \tilde{K}_2 \end{bmatrix} \begin{bmatrix} X \\ x \end{bmatrix}.$$
(11.2.16) ["simple15 "]

For the worst case shock chosen by the household to be aligned with the one chosen by the planner requires that

$$\tilde{K}_1 + \tilde{K}_2 = K.$$
 (11.2.17) ["simple16 "]

It can be proved that this equality holds, and furthermore that in this simple example that $\tilde{K}_2 = 0$, so that the household feeds back only on the part of the aggregate state that drives the scaled price process p_t .
As an example, set $\beta = .99, \theta = .001^{-1}, d_t - .7d_{t-1} = (1 - .7)5 + .05w_t, b_t = 30$. Let the state be $x_t = \begin{bmatrix} 1 \\ d_t \end{bmatrix}$. We computed that $\tilde{K}_1 = K = \begin{bmatrix} -0.0044 & 0.0001 \end{bmatrix}$ and $\tilde{K}_2 = 0.^1$

11.3. A robust planning problem

When the representative consumer fears model misspecification, the planning problem takes the form of a robust linear regulator problem. Let V(x) = -x'Px - p be the value of the robust planning problem starting from initial state x. The Bellman equation is

$$-x'Px - p = \min_{w} \max_{c,i,g} \left\{ -.5[(s-b) \cdot (s-b) + g \cdot g] + \beta \theta w'w + \beta E \left(-x^{*'}Px^{*} - p \right) \right\}$$
(11.3.1) ["rod13 "]

where the extremization is subject to the linear constraints

$$\begin{split} \Phi_c c + \Phi_g g + \Phi_i i &= \Gamma k + d & (11.3.2a) \quad ["rod14;a"] \\ k^* &= \Delta_k k + \Theta_k i & (11.3.2b) \quad ["rod14;b"] \\ h^* &= \Delta_h h + \Theta_h c & (11.3.2c) \quad ["rod14;c"] \\ s &= \Lambda h + \Pi c & (11.3.2d) \quad ["rod14;d"] \end{split}$$

$$z^* = A_{22}z + C_2(\epsilon + w)$$
 (11.3.2e) ["rod14;e"]

$$b = U_b z$$
 (11.3.2f) ["rod14;f"]

$$d = U_d z,$$
 (11.3.2g) ["rod14;g"]

where * denotes a next period value. Problem (11.3.1), (11.3.2) differs from (10.5.6), (10.5.7) in the following respects: (1) the addition of the distortion C_2w to the law of motion for z, (2) the appearance of $\beta\theta w'w$ in the continuation value function in (11.3.1), and (3) the minimization over w. As usual, $\theta > 0$ is a robustness parameter.

A Markov perfect equilibrium of the two-player zero-sum game (11.3.1), (11.3.2) is a pair of decision rules $u = -F(\theta)x, w = K(\theta)x$. The equilibrium determines two laws of motion for the state, namely,

$$x_{t+1} = A^o x_t + C \epsilon_{t+1}$$
 (11.3.3) ["nasset3 "]

¹ The Matlab program decentr.m computes an example.

and

$$x_{t+1} = (A^o + CK(\theta))x_t + C\epsilon_{t+1}, \qquad (11.3.4) \quad ["nasset4"]$$

where $A^o = A - BF(\theta)$. Equation (11.3.3) is the approximating model under the robust rule, while (11.3.4) is the worst case model under the robust rule. In chapter 12, we show that both of these models can be used to price assets by appropriately adjusting the stochastic discount factor.

11.4. Min-max representation of the household's problem in an Arrow-Debreu equilibrium

As in chapter 10, the value function and the robust law of motion for the state x_t contain information about competitive equilibrium prices (see formulas (10.5.14) and (10.5.21)). To decentralize the solution of the robust planning problem, we must verify that when the robust representative household faces those prices as a price-taker, he chooses the allocation that solves the planning problem. In checking that claim, we also verify that the representative household chooses a worst case model that is aligned with the planner's worst case model.

An Arrow-Debreu setting makes it natural to formulate the household's problem as a game in which the maximizing and minimizing players both commit to sequences, as in the sequence games of chapter 6. We find it convenient to embrace a recursive representation of each player's problem (see section 6.5.1). The maximizing player chooses an allocation, taking as given the law of motion for X chosen by the minimizing player. We use a guess-and-verify strategy. First we take up the problem of the maximizing player, having guessed that the minimizing player has chosen as the law of motion for X the same worst case law that is chosen by the robust planner. Second, setting the allocation equal to the one chosen by the planner, we verify that the minimizing player chooses the same distorted law of motion as does the planner.

11.4.1. Sequence problem of maximizing player

We first pose the problem of the maximizing player in a sequence formulation of the household's robust decision problem. As in chapter 10, we use a recursive representation of prices in the Arrow-Debreu equilibrium. This enables us to give a recursive representation of the maximizing player's decision problem. The maximizing player in the robust household chooses stochastic processes for $\{c_t, i_t, s_t, h_t, k_t, \ell_t\}_{t=0}^{\infty}$, each element of which is in L_0^2 , to maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t \left\{ -.5 \left[(s_t - b_t) \cdot (s_t - b_t) + \ell_t^2 \right] + \beta \theta X_{t+1}' K' K X_{t+1} \right\}$$
(11.4.1) ["rof1 "]

subject to

$$\begin{split} E_0 \sum_{t=0}^{\infty} \beta^t \left(p_{ct}^0 \cdot c_t + p_{it}^0 \cdot i_t \right) \mid J_0 \\ &= E_0 \sum_{t=0}^{\infty} \beta^t \left(p_{\ell t}^0 \ell_t + p_{rt}^0 \cdot k_{t-1} + p_{dt}^0 \cdot d_t \right) \mid J_0 \end{split}$$
(11.4.2*a*) ["rof2;a "]

$$\begin{split} s_t &= \Lambda h_{t-1} + \Pi c_t & (11.4.2b) \quad ["rof2;b"] \\ h_t &= \Delta_h h_{t-1} + \Theta_h c_t & (11.4.2c) \quad ["rof2;c"] \end{split}$$

$$k_t = \Delta_k k_{t-1} + \Theta_k i_t \tag{11.4.2d} ["rof2;d"]$$

$$b_t = U_b z_t$$
 (11.4.2e) ["rof2;e"]

$$d_t = U_d z_t$$
 (11.4.2*f*) ["rof2;f"]

$$p_t^0 = M X_t$$
 (11.4.2*h*) ["rof2;h"]

$$X_{t+1} = (A^o + CK)X_t + C\epsilon_{t+1}$$
(11.4.2*i*) ["rof2;i"]

Tom and Lars: do we want to include the law of motion for z_{t+1} here? with h_{-1}, k_{-1}, z_0 given, and where the mathematical expectation is evaluated with respect to the distribution over histories ϵ^t . In (11.4.2*i*), $K = K(\theta)$ is the matrix that defines the decision rule for the worst case shock in the robust planning problem. To compute the pricing matrix M in (11.4.2*h*), we use the counterparts to formulas (10.5.14), (10.5.21) in which A^o is the now the matrix A - BF that solves the robust planning problem and P is the matrix in the value function for the robust planning problem that appears on the left side of (11.3.1). In the household's problem (11.4.1), (11.4.2), E_0 denotes the expectation evaluated with respect to the ϵ_t 's that together with the distorted law of motion (11.4.2i) for the state X_{t+1} generate the distribution used to evaluate the present values in the budget constraint as well as the component of the state z_t that drives the preference shock b_t and the endowment process d_t . A convenient feature of problem (11.4.1), (11.4.2) is that the *same* distribution is used to evaluate future prospects both in the objective function (11.4.1) and in the budget constraint (11.4.2a). Putting a multiplier μ_0^w on the household's budget constraint, we can combine (11.4.1) and (11.4.2a) into a Lagrangian

The household's problem is to choose an allocation to maximize (11.4.3) and a multiplier μ_0^w to minimize it, subject to (11.4.2*b*)–(11.4.2*c*). For a given μ_0^w , this problem takes the form of an optimal linear regulator. In the following subsection 11.4.2, we give a formula that allows us to compute μ_0^w in advance directly from the allocation that solves the robust planning problem and the price system (11.4.2*h*), (11.4.2*i*). Being a linear regulator problem, it follows that the solution of the maximizing player's problem is a decision rule that can be expressed

$$\begin{bmatrix} c'_{t} & i'_{t} & \ell'_{t} \end{bmatrix}' = \overline{S} \begin{bmatrix} h_{t-1} \\ k_{t-1} \\ X_{t} \end{bmatrix}.$$
 (11.4.4) ["rof22 "]

For (11.4.4) to confirm the solution of the planning problem it must be true that

$$\overline{S} \begin{bmatrix} h_{t-1} \\ k_{t-1} \\ X_t \end{bmatrix} \equiv SX_t, \qquad (11.4.5) \quad ["rof24"]$$

when we set $h_{t-1} = H_{t-1}$ and $k_{t-1} = K_{t-1}$, where SX_t is the decision rule that solves the planning problem. Equality (11.4.5) assures that the robust household in the competitive equilibrium chooses the allocation that solves the planning problem.

The role of the pair (11.4.2h), (11.4.2i) is to provide a recursive representation of the prices that appear in the household's budget constraint (11.4.2a). In particular, notice how (11.4.2i) makes sure that in (11.4.2a) the conditional expectation is taken with respect to the 'twisted' or worst case distribution that emerges from the robust planning problem.² Equality (11.4.5) assures that the allocation that solves the planning problem satisfies the household's budget constraint (11.4.2a).

Another role of (11.4.2i) is to force the household to accept the twisted probability distribution chosen by the robust planner in evaluating conditional expectations of the shocks b_t, d_t . This streamlines the problem because it allows us to use a common distribution to express the conditional expectations E_0 that appear both in the objective function (11.4.1) and in the Harrison-Kreps representation of the budget constraint (11.4.2a).

11.4.2. Digression about computing μ_0^w

To solve problem (11.4.1), (11.4.2), it is useful to have a formula for the Lagrange multiplier μ_0^w on the household's budget constraint (11.4.2*a*) in a competitive equilibrium.³ Hansen and Sargent derive the following convenient formula for μ_0^w . They define the implicit price of consumption services as

$$\rho_t^0 \equiv \Pi^{-1\prime} \Big[p_{ct}^0 - \Theta_h' E_t \sum_{\tau=1}^\infty \beta^\tau (\Delta_h' - \Lambda' \Pi^{-1\prime} \Theta_h')^{\tau-1} \Lambda' \Pi^{-1\prime} p_{c,t+\tau}^0 \Big]. \tag{11.4.6} \quad \texttt{["oi8"]}$$

To compute μ_0^w , Hansen and Sargent (200XX) partition household capital and service sequences into two components. One is a service sequence obtained from the *initial* endowment of household capital h_{-1} . The other is the service sequence obtained from *market purchases* of consumption goods. The service sequence $\{s_{i,t}\}$ obtained from the initial endowment of household capital evolves according to:

$$s_{i,t} = \Lambda h_{i,t-1}$$
 (11.4.7) ["oi9"]
 $h_{i,t} = \Delta_h h_{i,t-1}$

where $h_{i,-1} = h_{-1}$. The service sequence $\{s_{m,t}\}$ obtained from purchases of consumption satisfies

$$s_{m,t} = b_t - s_{i,t} - \mu_0^w \rho_t^0.$$
(11.4.8) ["oi10"]

² Recall also that the A^o, P that are used in forming the prices via formulas (10.5.14), (10.5.21) pertain to the solution of the robust planning problem.

³ Counterparts of the same formulas would work for computing μ_0^w of the chapter 10 models without fear of model misspecification.

There are two ways to compute the time zero cost of the sequence $\{s_{m,t}\}$. One is to compute the time zero cost of the consumption sequence $\{c_t\}$ needed to support the service demands using the price sequence $\{p_{ct}^0\}$. Another is to use the implicit rental sequence $\{\rho_t^0\}$ directly to compute the time zero costs of $\{s_{m,t}\}$. Hansen and Sargent (200XXX) verify that the two measures of costs agree:

$$E_0 \sum_{t=0}^{\infty} \beta^t \rho_t^0 \cdot s_{m,t} = E_0 \sum_{t=0}^{\infty} \beta^t p_t^0 \cdot c_t.$$
(11.4.9) ["oil1"]

It follows from (11.4.8) that

$$E_0 \sum_{t=0}^{\infty} \beta^t \rho_t^0 \cdot s_{m,t} = E_0 \sum_{t=0}^{\infty} \beta^t \rho_t^0 \cdot (b_t - s_{i,t}) - \mu_0^w E_0 \sum_{t=0}^{\infty} \beta^t \rho_t^0 \cdot \rho_t^0. \quad (11.4.10) \quad ["\texttt{oil2"}]$$

Substitute (11.4.9) and (11.4.10) into the consumer's budget constraint (6.2), and solve for the time zero marginal utility of wealth μ_0^w :

$$\mu_0^w = \frac{E_0 \sum_{t=0}^\infty \beta^t \rho_t^0 \cdot (b_t - s_{i,t}) - W_0}{E_0 \sum_{t=0}^\infty \beta^t \rho_t^0 \cdot \rho_t^0}, \qquad (11.4.11) \quad ["oil3"]$$

where W_0 denotes initial period wealth given by

$$W_0 = E_0 \sum_{t=0}^{\infty} \beta^t (p_{\ell t}^0 \ell_t + p_{rt} \cdot k_{t-1} + p_{dt}^0 \cdot d_t).$$
(11.4.12) ["oi14"]

The geometric sums in (11.4.10) and (11.4.12) can be computed by solving Sylvester equations.

11.4.3. Sequence problem of minimizing player

Our next task is to pose the problem of the minimizing player in a sequence version of the zero-sum two-player game that describes the robust household. We freeze the allocation chosen by the maximizing player at the one described in the last section, which equals that chosen by a robust planner. As remarked in the previous subsection, that guarantees that the allocation satisfies the household's budget constraint (11.4.2*a*). We have to verify that the household chooses the same twisted law of motion that the robust planner did.

It simplifies things that the preceding argument allows us to drop the consumer's budget constraint because we know that it is satisfied at the allocation that solves the robust planning problem. The minimizing player within the household chooses a sequence $\{w_{t+1}\}_{t=0}^{\infty}$ to minimize

$$E_0 \sum_{t=0}^{\infty} \beta^t \left\{ -.5[(s_t - b_t) \cdot (s_t - b_t) + \ell_t^2] + \beta \theta w_{t+1} \cdot w_{t+1} \right\}$$
(11.4.13) ["mp101 "]

subject to

$$s_t = \Lambda h_{t-1} + \Pi c_t$$
 (11.4.14*a*) ["mp2;a "]

$$h_t = \Delta_h h_{t-1} + \Theta_h c_t$$
 (11.4.14b) ["mp2;b"]

$$k_t = \Delta_k k_{t-1} + \Theta_k i_t$$
 (11.4.14*c*) ["mp2;c "]

$$\begin{bmatrix} c'_t & i'_t & \ell'_t \end{bmatrix}' = S \dot{X}_t$$
(11.4.14d) ["mp2;d "]

$$X_{t+1} = (A^o + CK)X_t + C\epsilon_{t+1}$$
(11.4.14e) ["mp2;e "]

$$b_t = U_b z_t$$
 (11.4.14f) ["mp2;f"]

$$d_t = U_d z_t$$
 (11.4.14g) ["mp2;g "]

$$z_{t+1} = A_{22}z_t + C_2(\epsilon_{t+1} + w_{t+1}).$$
(11.4.14*h*) ["mp2;h "]

Here \tilde{X}_t is a state vector with components comparable to those in x_t , $A^o = A - BF(\theta)$ is the robust decision rule that the robust planner would choose, and we impose the initial condition $\tilde{X}_0 = x_0$. Also, the conditional expectation is evaluated with respect to the distribution of future values of the joint state h, k, z, \tilde{X} that is generated by the distribution for the ϵ_{t+1} 's together with the given laws of motion for the state variables. The role of (11.4.14*d*), (11.4.14*e*) is to freeze the household's allocation at the solution of the planning problem. The solution of the minimizing player's problem is a decision rule of the form

$$w_{t+1} = \tilde{K} \begin{bmatrix} X_t \\ x_t \end{bmatrix}. \tag{11.4.15} \quad ["mp3 "]$$

We want verify that

$$\tilde{K}\begin{bmatrix}\tilde{X}_t\\x_t\end{bmatrix} = K\tilde{X}_t \tag{11.4.16} \quad ["mp4 "]$$

when $x_t = \tilde{X}_t$. But problem (11.4.13), (11.4.14) is a recursive representation of the sequence version of the minimizing part of the robust planner's problem. Thus, (11.4.16) follows directly from results in chapter 6, section 6.5.1.

11.5. A decentralization with Arrow securities

11.5.1. A robust consumer trading Arrow securities

This section adapts the equilibrium with sequential trading of one-period Arrow securities, as in section 10.6, to the situation where the representative household is concerned about model misspecification.

As a price taker, the household faces the one-step ahead pricing kernel $q(X_{t+1}, X_t)$ that obeys the following version of (10.6.4)

$$q(X_{t+1}|X_t) = \beta \frac{e_1 u_{c,t+1}}{e_1 u_{c,t}} \hat{f}(X_{t+1}|X_t)$$
(11.5.1*a*) ["arrow6R;a"]

where the conditional density $\hat{f}(X_{t+1}|X_t)$ is induced by the difference equation

$$X_{t+1} = (A^o + CK)X_t + C\epsilon_{t+1}$$
(11.5.1b) ["arrow6R;b"]

where $\epsilon_{t+1} \sim \mathcal{N}(0, I)$, $A^o = A - BF(\theta)$, and where $A^o + CK$ is the worst-case transition matrix for the autonomous law of motion for x_t that emerges from the robust planning problem for a given θ .⁴

We obtain the household's Bellman equation by replacing (10.6.5), (10.6.6) with a two-player, zero-sum game. We generate a recursion in a value function $W(a_t, h_{t-1}, k_{t-1}, X_t)$ by posing an 'inner problem' by which a maximizing player chooses $c_t, i_t, \ell_t, a(X_t)$, taking as given the feedback rule for w_{t+1} ; and an 'outer

⁴ Notice that $\hat{f}(X_{t+1}|X_t)$ is the transition density under the planner's worst case model and the robust rule. See chapter 12 for representations of pricing functions in terms of the approximating model.

problem' in which a minimizing player who chooses w_{t+1} takes as given the decision rule for $c_t, i_t, \ell_t, a(X_t)$ chosen by the maximizing player.

11.5.2. The inner problem

The maximizing player solves

$$\max_{c_{t},i_{t},\ell_{t},a(X_{t+1})} \left\{ -(|s_{t} - b_{t}|^{2} + \ell_{t}^{2}) + \beta \theta w_{t+1}' w_{t+1} + \beta E_{t} \int W(a(X_{t+1}),h_{t},k_{t},X_{t+1}) \right\}$$
(11.5.2) ["arrow10R "]

where the maximization is subject to

$$\begin{split} s_t &= \Lambda h_t + \Pi c_t & (11.5.3a) \quad ["arrow11R;a"] \\ h_t &= \Delta_h h_{t-1} + \Theta_h c_t & (11.5.3b) \quad ["arrow11R;b"] \\ k_t &= \Delta_k k_{t-1} + \Theta_k i_t & (11.5.3c) \quad ["arrow11R;c"] \\ X_{t+1} &= A^o X_t + C(\epsilon_{t+1} + w_{t+1}) & (11.5.3d) \quad ["arrow11R;d"] \\ w_{t+1} &= K X_t & (11.5.3e) \quad ["arrow11R;c"] \\ p_t &= M X_t & (11.5.3f) \quad ["arrow11R;f"] \\ b_t &= S_b X_t & (11.5.3f) \quad ["arrow11R;f"] \\ d_t &= S_d X_t & (11.5.3h) \quad ["arrow11R;f"] \\ \int a(X_{t+1})q(X_{t+1}|X_t)dX_{t+1} &= a_t + p_{\ell t}\ell_t + p_{rt} \cdot k_{t-1} \\ &+ p_{dt} \cdot d_t - p_{ct} \cdot c_t - p_{it} \cdot i_t. & (11.5.3i) \quad ["arrow11R;i"] \end{split}$$

Here E_t denotes the mathematical expectation generated with respect to (11.5.3*d*), (11.5.3*e*). In (11.5.2), (11.5.3), the maximizing player takes for granted that w_{t+1} conforms to the decision rule chosen by the robust planner, namely,

$$w_{t+1} = KX_t \qquad (11.5.4) \quad ["wrobplanner"]$$

and chooses a decision rule of the form

$$\begin{bmatrix} c_t & i_t & \ell_t & a(X_t) \end{bmatrix} = S \begin{bmatrix} h_{t-1} \\ k_{t-1} \\ X_t \end{bmatrix}.$$
 (11.5.5) ["arrow12R "]

11.5.3. The outer problem

The problem of the minimizing player is

$$W(a_t, h_{t-1}, X_t) = \min_{w_{t+1}} \left\{ -(|s_t - b_t|^2 + \ell_t^2) + \beta \theta w'_{t+1} w_{t+1} + \beta E_t \int W(a(X_{t+1}), h_t, k_t, X_{t+1}) \right\}$$
(11.5.6) ["arrow10R "]

where the maximization is subject to equation (11.5.5) and

$$s_t = \Lambda h_t + \Pi c_t \qquad (11.5.7a) \quad ["arrow11R;a"]$$

$$h_t = \Delta_h h_{t-1} + \Theta_h c_t \qquad (11.5.7b) \quad \texttt{["arrow11R;b"]}$$

$$k_t = \Delta_k k_{t-1} + \Theta_k i_t \qquad (11.5.7c) \quad ["arrow11R;c"]$$

$$X_{t+1} = A^o X_t + C(\epsilon_{t+1} + w_{t+1})$$
(11.5.7d) ["arrow11R;d"]

$$b_t = S_b X_t$$
 (11.5.7*f*) ["arrow11R;f"]

The minimizing player also takes as given the decision rule (11.5.5) emerging from the inner problem and chooses a decision rule for w_{t+1} of the form

$$w_{t+1} = \tilde{K} \begin{bmatrix} h_{t-1} \\ k_{t-1} \\ X_t \end{bmatrix}.$$

We have not included the budget constraint (11.5.3*h*) in the outer problem because we know that it is satisfied from the way the inner problem has constructed decision rule (11.5.5). The optimal decision rule of the minimizing player has the form $w_{t+1} = \tilde{K} \begin{bmatrix} h_{t-1} \\ k_{t-1} \\ X_t \end{bmatrix}$. This structure can be shown to affirm the identity $\tilde{K} \begin{bmatrix} h_{t-1} \\ k_{t-1} \\ X_t \end{bmatrix} \equiv KX_t$, which aligns the worst case shock in a recursive competitive equilibrium with that emerging from a robust planning problem.

11.6. An ex post Bayesian planning problem

There is also an expost Bayesian version of the planning problem like the one described in chapter 6. It endows the planner with a belief about the law of motion that is distorted relative to his approximating model in just such a way that he attains a robust rule by solving an ordinary Bellman equation without a concern for robustness.

We can apply an idea of chapter 6 to confront the planner with a distorted law of motion for z that will inspire him to choose a robust decision rule. This leads to what is known as an expost Bayesian problem because the robust rule cannot be dominated in the sense of statistical decision theory. Rather than assigning the planner the approximating model, we endow him with a model that has been twisted to promote robustness. Taking that model as given, the planner then behaves as an ordinary planner without concern about misspecification.

To formulate this expost Bayesian problem, we can augment the state variable x by a vector X of the same dimension. The Bellman equation is

$$-x'Px - p = \max_{c,i,g} \left\{ -.5[(s-b) \cdot (s-b) + g \cdot g] + \beta E \left(-x^{*'}Px^{*} - p \right) \right\}$$
(11.6.1) ["rod15 "]

subject to the linear constraints formed by (11.3.2) and the following additional exogenous law of motion for w:

$$X^* = A^o X + C(\epsilon + w)$$
(11.6.2*a*) ["rod16;a"]

$$w = KX.$$
 (11.6.2b) ["rod16;b"]

Three features are noteworthy relative to equations (11.3.2). First, (11.6.1) is an ordinary (non-robust) dynamic programming problem. Second, through equation (11.3.2e), the w determined by equation (11.6.2) feeds back on the z process that governs the shocks impinging on the consumer's preference and endowment shock processes, (b, d). Third, we have augmented the state by 'big X' in order to have the ex post Bayesian planner take w as exogenous and beyond his control. Where u_t is the control, the planner chooses a decision rule $u = -\tilde{F} \begin{bmatrix} x \\ X \end{bmatrix}$. Chapter 6 shows how, after equating X = x, this decision rule satisfies

$$-\tilde{F}\begin{bmatrix}x\\x\end{bmatrix} = -Fx$$

where u = -Fx solves the robust planning problem (11.3.1), (11.3.2).

11.6.1. Remarks on practicality

The direct way of solving the robust planning problem is obviously the more useful one computationally, not only because it has a lower dimensional state vector (x as compared with (x, X)), but also because in order to find the distorted model (11.6.2), we have first to solve a robust planning problem. Nevertheless, the *ex post* Bayesian method is a useful reinterpretation of the allocation associated with the robust planning problem, one that we shall use when we turn to asset pricing.

11.7. Two asset pricing strategies

Under a preference for robustness, there are two counterparts to the strategy described in section 10.7 that express asset prices in terms of the state vector x_t of the planner. The first uses the robust planning problem (11.3.1), (11.3.2), and the second uses the *ex post* Bayesian planning problem (11.6.1), (11.6.2).

11.7.1. Pricing from the robust planning problem

Hansen, Sargent, and Tallarini (1999) used the following three step method to compute asset prices:

- 1. Solve the robust planning problem.
- 2. Obtain representations for the planner's shadow prices, based on marginal utilities of consumption evaluated at the allocation that solves the planning problem.
- 3. Use the appropriate shadow prices to price assets as conditional expectations of inner products of (scaled) history-date prices, computing the conditional expectation by taking the distorted law of motion cast in terms of little x that emerges from the robust planning problem and the corresponding sequence of information J_t .

This method leads to a representation for asset prices corresponding to (10.7.1) of the following form:

$$x_{t+1} = \hat{A}^o x_t + C \epsilon_{t+1}$$
 (11.7.1*a*) ["nutshell2;a"]

$$y_t = S_y x_t \qquad (11.7.1b) \quad ["nutshell2;b"]$$

$$p_{cs}^t = (e_1 M_c x_t)^{-1} M_c x_s$$
 (11.7.1c) ["nutshell2;c"]

$$a_{yt} = E_t \sum_{t=0}^{\infty} \beta^t p_{ct+j}^t \cdot y_{t+j}, \qquad (11.7.1d) \quad ["nutshell2;d"]$$

where E is the expectation evaluated with respect to $\hat{A}^o = A^o + CK$, the transition matrix for the worst-case transition law under the robust decision rule F, where $A^o = A - BF$; and M_c also incorporates the worst case transition law through the presence of $A^o + CK$ in (10.5.12), (10.5.11) (also see formulas (10.5.14) and (10.5.21)).

11.7.2. Pricing from the expost Bayesian planning problem

An alternative strategy is based on the following four step procedure:

- 1. Solve a robust planning problem.
- 2. Obtain a representation of the worst case shock process in terms of the new state variable X_t as on page 263.
- 3. Solve the ordinary (i.e., non-robust) planning problem with the distorted law of motion for the augmented state $\begin{bmatrix} x_t \\ X_t \end{bmatrix}$ as in (11.6.1), (11.6.2).
- 4. Use either of the complete-market decentralizations presented in chapter 10 for our economies *without* robustness and price assets using the standard (non-robust) asset pricing formulas.

11.8. Concluding remarks

Chapter 10 showed that without a preference for robustness, the pricing kernel for Arrow securities has the representation

$$q(X_{t+1}|X_t) = \beta \frac{e_1 u_{c,t+1}}{e_1 u_{c,t}} f(X_{t+1}|X_t)$$

where $f(X_{t+1}|X_t)$ is the transition density under the approximating model. In this chapter, we have shown that a representative consumer's fear that the approximating model is misspecified makes the pricing kernel for one-period Arrow securities become

$$q(X_{t+1}|X_t) = \beta \frac{e_1 u_{c,t+1}}{e_1 u_{c,t}} \hat{f}(X_{t+1}|X_t)$$

where $\hat{f}(X_{t+1}|X_t)$ is the planner's worst-case model. The price of Arrow securities under a preference for robustness can also be written as

$$q(X_{t+1}|X_t) = \beta \frac{e_1 u_{c,t+1}}{e_1 u_{c,t}} \left(\frac{\hat{f}(X_{t+1}|X_t)}{f(X_{t+1}|X_t)} \right) f(X_{t+1}|X_t).$$
(11.8.1) ["keyassetformula "]

In chapter 12, we shall use representations like (11.8.1) to price assets under the approximating model. The term $\left(\frac{\hat{f}(X_{t+1}|X_t)}{f(X_{t+1}|X_t)}\right)$ can be viewed as a multiplicative adjustment to the usual stochastic discount factor $\beta \frac{e_1 u_{c,t+1}}{e_1 u_{c,t}}$ that is contributed by the representative household's concern about model misspecification.⁵ When the likelihood ratio $\left(\frac{\hat{f}(X_{t+1}|X_t)}{f(X_{t+1}|X_t)}\right)$ is volatile under the approximating model, it serves to boost the market price of macroeconomic risk.

⁵ The likelihood ratio $(\frac{f}{f})$ that adjusts the stochastic discount factor for robustness also governs the detection error probability statistics described in chapter 8. See chapter 12 and Anderson, Hansen, and Sargent (2003) for more about the connection between detection error statistics and theoretical values of market prices of risk.

A. Decentralization of partial equilibrium

For a partial equilibrium model with adjustment costs, this appendix studies a recursive competitive equilibrium in which the representative firm has a preference for robust decisions. We show that the standard trick of computing an equilibrium by solving the fictitious planning problem of maximizing a discounted sum of consumer plus producer surplus extends to a setting where the firm wants robustness. In this case, the planning problem becomes a robust planning problem in which the planner extremizes over decision, model distortion pairs.

Consider an adaptation for robustness of Sargent's (1987, chapter XVI) version of Lucas and Prescott's model of investment under uncertainty. Demand for a single good is governed by an inverse demand function

$$p_t = A_0 - A_1 \overline{q}_t + v_t$$
 (11.A.1) ["lpr1 "]

where

$$v_{t+1} = \rho v_t + C_v w_{t+1}.$$
 (11.A.2) ["lpr2 "]

A representative firm has one-period quadratic cost function $\sigma(q_t, q_{t+1})$ and one-period profits $\pi_t = p_t q_t - \sigma(q_t, q_{t+1})$. The firm acts as a price taker and wants to extremize $\sum_{t=0}^{\infty} \beta^t(p_t q_t - \sigma(q_t, q_{t+1}))$ with respect to sequences for $\{q_{t+1}, w_{t+1}\}_{t=0}^{\infty}$. The firm believes that the law of motion for aggregate output is

$$\overline{q}_{t+1} = \ell_q(\overline{q}_t, v_t) \tag{11.A.3} ["lpr3 "]$$

where ℓ_q is a linear function. The representative firm solves the two-player zero-sum game

$$\min_{\{w_{t+1}\}} \max_{\{q_{t+1}\}} \sum_{t=0}^{\infty} \beta^t \{ p_t q_t - \sigma(q_t, q_{t+1}) + \beta \theta w_{t+1}^2 \}$$
(11.A.4) ["lpr3a "]

where the extremization is subject to (11.A.1), (11.A.2), (11.A.3). An equilibrium of the representative agent's two-player zero-sum game is a pair of decision rules

$$q_{t+1} = \phi_q(q_t, v_t, \overline{q}_t)$$
 (11.A.5a) ["oli111;a"]

$$w_{t+1} = \phi_w(q_t, v_t, \overline{q}_t).$$
 (11.A.5b) ["oli111;b"]

The representative agent's extremization problem induces a mapping from ℓ_q in (11.A.3) to (ϕ_q, ϕ_w) . When the representative firm perceives the law of motion for \overline{q}_t to be (11.A.3), it acts to make the actual law of motion to be $\overline{q}_{t+1} = \phi_q(\overline{q}_t, v_t, \overline{q}_t)$. A competitive equilibrium under robustness is a fixed point of the mapping from $\ell_q(\overline{q}, v)$ to $\phi_q(\overline{q}, v, \overline{q})$. That is, for the representative firm to be representative, it must be true that ℓ_q satisfies

$$\phi_q(\overline{q}_t, v_t, \overline{q}_t) = \ell_q(\overline{q}_t, v_t). \tag{11.A.6} ["lpr5 "]$$

Fortunately, by extending lines of argument of Lucas and Prescott (1972) and Sargent (1987), it is not necessary to attack this fixed point problem directly. In particular, we can compute ℓ_q and an associated ℓ_w directly by solving a fictitious robust planning problem. The fictitious planning problem is

$$\min_{\{w_{t+1}\}} \max_{\{\overline{q}_{t+1}\}} \sum_{t=0}^{\infty} \beta^t \{ S(\overline{q}_t, v_t) - \sigma(\overline{q}_t, \overline{q}_{t+1}) + \beta \theta w'_{t+1} w_{t+1} \}$$
(11.A.7) ["lpr6 "]

where $S(\overline{q}, v)$ is consumer surplus defined as

$$S(\overline{q}, v) = \int_0^{\overline{q}} (A_0 - A_1 x + v) dx$$
$$= A_0 \overline{q} - \frac{A_1}{2} \overline{q}^2 + \overline{q} v.$$

The state of the market is \overline{q}_t, v_t . A solution of this two-player zero-sum game is

$$\overline{q}_{t+1} = \ell_q(\overline{q}_t, v_t) \tag{11.A.8a} \quad \texttt{["partial1;a"]}$$

$$w_{t+1} = \ell_w(\overline{q}_t, v_t) \tag{11.A.8b} ["partial1;b"]$$

It turns out that $\ell_j(\overline{q}, v) = \phi_j(\overline{q}, v, \overline{q})$ for j = q, w. This assertion can be proved by extending the proof in Sargent (1987, ch. XIV). The proof strategy is to obtain the Euler equations for extremizing (11.A.4), then to use the demand curve (11.A.1) to eliminate price, rearrange, and note that these Euler equations-cum equilibrium conditions match the Euler equations for extremizing the fictitious planning criterion (11.A.7).

Chapter 12 Asset pricing

12.1. Introduction

This chapter explores how a fear of model misspecification affects prices of risky securities. Without fear of misspecification, the price of a claim to a random future payoff equals the conditional expectation of the inner product of a stochastic discount factor and the random future payoff, evaluated using the representative agent's model.¹ When the representative agent fears misspecification of his approximating model, two such inner-product representations of asset prices are available. They differ in what they take as the model with respect to which the conditional expectation is evaluated. In one, the conditional expectation is evaluated with respect to the representative agent's worst case model, a model that depends on the parameter θ that calibrates his fear of misspecification. A second representation of the same prices exists because the approximating model and the worst case model put positive probabilities on the same events. This second representation evaluates the conditional expectation with respect to the approximating model. The first representation captures a concern about robustness by adjusting the probability distribution relative to the approximating model, while the second representation instead adjusts the stochastic discount factor (a.k.a. pricing kernel). In particular, to represent asset prices in terms of conditional expectations under the approximating model, the second representation multiplies the ordinary stochastic discount factor without fear of misspecification by the likelihood ratio, or Radon-Nikodym derivative, of the endogenous worst case distorted model relative to the approximating model. The expected value of that likelihood ratio is the entropy measure that we used in chapter 2 to measure the proximity of models. It also governs the detection statistics of chapter 8.

After reviewing asset pricing formulas in a standard model without a fear of misspecification, this chapter modifies those formulas to express a representative agent's fear of misspecification. By way of examples, we study asset pricing

 $^{^{1}}$ Without fear about misspecification, an agent can discard the adjective 'approximating'.

^{- 293 -}

in the permanent income economy of chapter 9 and a partial equilibrium occupational choice model of Jaewoo Ryoo and Sherwin Rosen (2003).

12.2. Approximating and distorted models

Chapters 10 and 11 describe planning problems and competitive equilibria for a class of linear-quadratic models of which the consumption smoothing model of chapter 9 and the occupational choice model of section 12.6 are special cases. The environment of chapter 10 is arranged so that without a fear of misspecification, the planning problem fits into the optimal linear regulator problem. Chapter 11 then uses a robust linear regulator to create a model in which the representative household's fear of misspecification is indexed by parameter $\theta > 0$. Equilibrium representations for prices and quantities can be determined from the solution of the robust linear regulator.

Chapter 10 describes matrices that portray the preferences, technology, and information structure of the economy. These can be assembled into matrices that define the robust linear regulator for a planning problem. The solution of the planning problem determines competitive equilibrium prices and quantities. Associated with the robust planning problem is the Bellman equation

$$-x'Px - p = \max_{u} \min_{w} \left\{ r(x, u) + \theta \beta w'w + \beta E(-x^{*'}Px^{*} - p) \right\}$$
(12.2.1)

where the extremization is subject to

$$x^* = Ax + Bu + C(\epsilon + w),$$
(12.2.2)

where $\epsilon \sim \mathcal{N}(0, I)$ and $\theta \in (\underline{\theta}, +\infty]$. A Markov perfect equilibrium of this two-player zero-sum game is a pair of decision rules $u = -F(\theta)x, w = K(\theta)x$. The equilibrium determines the following two laws of motion for the state that interest us:

$$x_{t+1} = A^o x_t + C\epsilon_{t+1} \tag{12.2.3}$$

and

$$x_{t+1} = (A^o + CK(\theta))x_t + C\epsilon_{t+1}, \qquad (12.2.4)$$

where $A^o = A - BF(\theta)$. For a given $\theta \in [\underline{\theta}, \infty)$, (11.3.3) is the approximating model under the robust rule for u, while (12.2.4) is the distorted worst case model under the robust rule.

Where there is no fear of misspecification, $\theta = +\infty$. Chapter 10 describes a class of economies whose equilibria can be presented in the form (12.2.4) together with selector matrices that determine equilibrium prices and quantities as functions of the state x_t . In particular, quantities Q_t and scaled statecontingent prices p_t are linear functions of the state:

$$Q_t = S_Q x_t \tag{12.2.5a}$$

$$p_t = p_Q x_t. \tag{12.2.5b}$$

We shall soon remind the reader what we mean by 'scaled prices'. We showed how to compute these in chapter 10; see formulas (10.5.14), (10.5.21).

To get equilibria under a fear of misspecification, we simply set $\theta < +\infty$ in (12.2.1). Formulas for equilibrium prices and quantities from chapter 10 (i.e., the S_Q, M_Q in (12.2.5)) apply directly. Associated with an equilibrium under a fear of misspecification are the approximating transition law (12.2.3) and the distorted transition law (12.2.4) for the state x_t , as well as auxiliary equations for prices and quantities of the form (12.2.5).

The approximating and distorted equilibrium laws of motion (12.2.3) and (12.2.4) induce Gaussian transition densities²

$$f(x_{t+1}|x_t) \sim \mathcal{N}(A^o x_t, CC') \tag{12.2.6a}$$

$$\hat{f}(x_{t+1}|x_t) \sim \mathcal{N}((A^o + CK)x_t, CC'),$$
 (12.2.6b)

where we use f without a $(\hat{\cdot})$ to denote a transition density under the approximating model and f with a $(\hat{\cdot})$ to denote a probability associated with the distorted model (12.2.4). These transition densities induce joint densities $f^{(t)}(x^t)$ on histories $x^t = [x_t, x_{t-1}, \ldots, x_0]$ via

$$f^{(t)}(x^{t}) = f(x_{t}|x_{t-1})f(x_{t-1}|x_{t-2})\dots f(x_{1}|x_{0})f(x_{0}),$$

and similarly for $\hat{f}^{(t)}(x^t)$. Let $f_t(x_t|x_0)$ denote the t-step transition densities

$$f_t(x_t|x_0) \sim \mathcal{N}(A^{ot}x_0, V_t) \tag{12.2.7a}$$

$$\hat{f}_t(x_t|x_0) \sim \mathcal{N}((A^o + CK)^t x_0, \hat{V}_t),$$
 (12.2.7b)

² An alternative formulation on page 399 allows a broader set of perturbations of a Gaussian approximating model by letting the minimizing agent choose an arbitrary density. Under that formulation, the minimizing agent would still choose a Gaussian transition density with the same conditional mean as (12.2.6b) but with conditional covariance $\hat{C}\hat{C}' = C(I - \theta^{-1}C'PC)^{-1}C'$.

where V_t satisfies the recursion $V_t = A^{o'}V_{t-1}A_o + CC'$ initialized from $V_1 = CC'$, and \hat{V}_t satisfies the recursion $\hat{V}_t = (A^o + CK)'\hat{V}_{t-1}(A_o + CK) + CC'$ initialized from $\hat{V}_1 = CC'$.

12.3. Asset pricing without robustness

In section 10.7, we explained how the value of claims on risky streams of returns can be represented as the inner product of price and payout processes, where both the price and payout are expressed as functions of the planner's state vector x_t . In portraying the household's problem in a recursive competitive equilibrium, we needed to distinguish between the individual household's x_t and its 'market wide' counterpart X_t that drives prices. Nevertheless, we showed that for the purpose of computing asset prices, we can exclude X_t from the state vector and simply use x_t as the state vector. Accordingly, in the remainder of this chapter, we express prices in terms of x_t and histories x^t .³

When $\theta = +\infty$, there is no discrepancy between the distorted and worst case models and the following standard representative agent asset pricing theory applies. Let c_t denote a vector of time-t consumption goods. The price of a unit vector of consumption goods in period t contingent on the history x^t is⁴

$$q^{(t)}(x^t|x_0) = \beta^t \frac{u'(c_t(x^t))}{e_1 \cdot u'(c_0(x_0))} f^{(t)}(x^t|x_0), \qquad (12.3.1)$$

where $c_t(x^t)$ is a possibly history-dependent state-contingent consumption process, u'(c) is the vector of marginal utilities of consumption, and e_1 is a selector vector that pulls off the first consumption good, the time-zero value of which we take as numeraire. To make (12.3.1) well defined, we assume that $e_1 \cdot u'_1(c_0(x_0)) \neq 0$ with probability one. If we assume that the consumption allocation is not history-dependent, so that $c_t(x^t) = c(x_t)$, as is true in the models that occupy us, then we can use the *t*-step pricing kernel

$$q_t(x_t|x_0) = \beta^t \frac{u'(c(x_t))}{e_1 \cdot u'(c(x_0))} f_t(x_t|x_0).$$
(12.3.2)

 $^{^3~}$ The household in a competitive economy would face prices that are the same functions of $X_t~$ and X^t .

⁴ We denote by $u'(c_t)$ the vector of marginal utilities of the consumption vector c_t . In our model, $u'(c_t) = M_c x_t$.

Let an asset entitle its owner to $\{y(x_t)\}_{t=0}^{\infty}$, a stream of a vector of consumption goods whose state-contingent price is given by (12.3.2). The time-0 price of the asset is

$$a_0 = \sum_{t=0}^{\infty} \int_{x_t} q_t(x_t | x_0) \cdot y(x_t) dx_t$$

or

$$a_0 = \sum_{t=0}^{\infty} \int_{x_t} \beta^t \frac{u'(c(x_t))}{e_1 \cdot u'(c(x_0))} y(x_t) f_t(x_t | x_0) dx_t.$$
(12.3.3)

We can represent (12.3.3) as

$$a_0 = \frac{E_0 \sum_{t=0}^{\infty} \beta^t u'(c(x_t)) \cdot y(x_t)}{e_1 \cdot u'(c(x_0))}.$$
(12.3.4)

In linear-quadratic general equilibrium models, $u'(c(x_t))$ and $y(x_t)$ are both linear functions of the state. This means that the price of an asset is the conditional expectation of a geometric sum of a quadratic form, as portrayed in (12.3.4). Equation (12.3.4) implies a Sylvester equation (see page 93). Thus, let

$$p_c(x_t) = \frac{u'(c(x_t))}{e_1 \cdot u'(c(x_0))}.$$

Then the asset price can be represented

$$a_0 = E_0 \sum_{t=0}^{\infty} \beta^t p_c(x_t) \cdot y(x_t).$$
 (12.3.5)

We can regard p_c as a scaled Arrow Debreu price: it equals the Arrow-Debreu state price divided by β^t times a conditional probability. Scaling the price system in this way facilitates computation of asset prices as conditional expectations of an inner product of state prices and pay outs. Often $\beta^t p_c(x_t)$ is called a *t*-period stochastic discount factor. Below we shall also denote the stochastic discount factor as $m_{0,t} \equiv \beta^t p_c(x_t)$, so that (12.3.5) becomes

$$a_0 = E_0 \sum_{t=0}^{\infty} m_{0,t} \cdot y(x_t).$$

Hansen and Sargent (200XX) provide a complete treatment of asset pricing within linear-quadratic general equilibrium models. They show that: (1) equilibrium scaled Arrow-Debreu prices and quantities have representations (12.2.5); (2) the information required to form the matrix S_Q is embedded in F, A, Bfrom the optimal linear regulator problem; and (3) the matrices M_p that pin down the scaled Arrow-Debreu prices can be extracted from the matrix P in the value function -x'Px - p and the matrix $A^o = A - BF$ that emerge from the planner's problem (see formulas (10.5.14), (10.5.21)). Thus, in such models

$$p_c(x_t) = M_c x_t / e_1 M_c x_0. (12.3.6)$$

See formulas (10.5.11), (10.5.13) in chapter 10 for a formula for M_c and more details.

12.4. Asset pricing with robustness

We activate a fear of misspecification by setting $\theta < +\infty$, which causes the transition densities (12.2.6*a*), (12.2.6*b*) under the approximating and distorted models to disagree. In addition, the formulas for S_Q and M_Q in (12.2.5) respond to the setting for θ , via the dependence of S_Q on $F(\theta)$ and the dependence of M_Q on the *P* that solves the Bellman equation (12.2.1). Again, see (10.5.14), (10.5.21). We give an example in section 12.6.

The price system that supports a competitive equilibrium can be represented in the forms (12.3.1) and (12.3.2), with the distorted densities $\hat{f}^{(t)}$ and \hat{f}_t replacing the corresponding densities for the approximating model in (12.3.1) and (12.3.2). Thus, with a fear of misspecification, the time 0 price of the asset corresponding to (12.3.3) is

$$a_0 = \sum_{t=0}^{\infty} \int_{x_t} \beta^t p_c(x_t) \cdot y(x_t) \hat{f}_t(x_t | x_0) dx_t.$$
(12.4.1)

We can represent (12.4.1) as

$$a_0 = \hat{E}_0 \sum_{t=0}^{\infty} \beta^t p_c(x_t) \cdot y_t$$
 (12.4.2)

where \hat{E} denotes mathematical expectation using the distorted model (12.2.4), and $u'(c(x_t))$ must be computed using the M_Q in representation (12.2.5b) associated with θ .

12.4.1. Adjustment of stochastic discount factor for fear of model misspecification

Formula (12.4.2) represents the asset price in terms of the distorted measure that the planner uses to evaluate future utilities in the Bellman equation (12.2.1). To compute asset prices using this formula, we must solve a Sylvester equation using transition matrix $A^o + CK(\theta)$ from equation (12.2.4) to reflect that we are evaluating the expectation using the *distorted* transition law. We can also evaluate asset prices by computing expectations under the *approximating* model, but this requires that we adjust the stochastic discount factor to make the asset price satisfy (12.4.1). By dividing and multiplying by $f_t(x_t|x_0)$, we can represent (12.4.1) as

$$a_0 = \sum_{t=0}^{\infty} \int_{x_t} \beta^t p_c(x_t) \left(\frac{\hat{f}_t(x_t|x_0)}{f_t(x_t|x_0)} \right) \cdot y(x_t) f_t(x_t|x_0) dx_t.$$
(12.4.3)

or

$$a_0 = E_0 \sum_{t=0}^{\infty} \beta^t p_c(x_t) \left(\frac{\hat{f}_t(x_t|x_0)}{f_t(x_t|x_0)} \right) \cdot y(x_t),$$
(12.4.4)

where the absence of a $(\hat{\cdot})$ from *E* denotes that the expectation is evaluated with respect to the approximating model (12.2.3).⁵

In summary, with a fear of misspecification, if we want to evaluate asset prices under the approximating model, we have to adjust the ordinary *t*-period stochastic discount factor $m_{0,t} = \beta^t p_c(x_t)$ for a concern about model misspecification and use the modified stochastic discount factor:

$$m_{0,t}\left(\frac{\hat{f}_t(x_t|x_0)}{f_t(x_t|x_0)}\right).$$

Such a multiplicative adjustment to the stochastic discount factor $m_{0,t}$ carries over to nonlinear models. For our linear-quadratic-Gaussian setting, the likelihood ratio is

$$L_t = \frac{\hat{f}_t(x_t|x_0)}{f_t(x_t|x_0)} = \exp\left[\sum_{s=1}^t \left\{\epsilon'_s w_s - .5w'_s w_s\right\}\right].$$

 $^{^{5}}$ Notice the appearance of the same likelihood ratio in (12.4.4) used to define entropy in chapters 2 and 17 and to describe detection error probabilities in chapter 8.

12.4.2. Reopening markets

This section describes how to extend our asset pricing formulas to allow us to price 'tail assets' that are traded at time t and that pay vectors of consumption $\{y_{\tau}\}_{\tau=t}^{\infty}$ for t > 0. We want the price to be stated in time t units of the numeraire good.

Letting the *t*-step discount factor at time 0 be $m_{0,t} \equiv \beta^t p_c(x_t)$, (12.4.2) can be portrayed as

$$a_0 = \hat{E}_0 \sum_{t=0}^{\infty} m_{0,t} \cdot y_t \tag{12.4.5}$$

where $m_{0,t}$ is a vector of time-0 stochastic discount factors for pricing a vector of time-t payoffs. Define $m_{t,\tau}$ as the vector of corresponding time-t stochastic discount factors for pricing time- $\tau \geq t$ payoffs:⁶

$$m_{t,\tau} = \beta^{\tau-t} p_c(x_\tau) / e_1 p_c(x_t).$$
(12.4.6)

Then in time t units of the numeraire consumption good, the vector of payoffs $\{y_{\tau}\}_{\tau=0}^{\infty}$ is

$$a_t = \hat{E} \sum_{\tau=0}^{\infty} m_{t,\tau} y_{\tau}.$$
 (12.4.7)

Equation (12.4.7) is equivalent with

$$a_t = E_t \sum_{\tau=0}^{\infty} (m_{t,\tau} m_{t,\tau}^u) \cdot y_{\tau}, \qquad (12.4.8)$$

where the appropriate multiplicative adjustment $m_{t,\tau}^u$ to the stochastic discount factor is the likelihood ratio

$$m_{t,\tau}^{u} = \frac{f_{\tau-t}(x_{\tau}|x_{t})}{f_{\tau-t}(x_{\tau}|x_{t})} = \exp\left[\sum_{s=t}^{\tau} \{\epsilon_{s}'w_{s} - .5w_{s}'w_{s}\}\right].$$
(12.4.9)

⁶ We assume that $e_1 p_c(x_t) \neq 0$ with probability 1.

12.5. Pricing single period payoffs

For the purpose of using the permanent income model of chapter 9 to shed light on the implications of a fear of misspecification for the equity premium, let consumption be a scalar process and y_{t+1} a scalar random payoff at time t+1. Without a fear of misspecification, the price at time t of a time t+1 payout is

$$a_t = E_t m_{t,t+1} y_{t+1}. (12.5.1)$$

Applying the definition of a conditional covariance to (12.5.1) and using the Cauchy-Schwartz inequality implies

$$\left(\frac{a_t}{E_t m_{t,t+1}}\right) \ge E_t y_{t+1} - \left(\frac{\sigma_t(m_{t,t+1})}{E_t m_{t,t+1}}\right) \sigma_t(y_{t+1}).$$
(12.5.2)

The bound is attained by payoffs on the efficient frontier. The left side is the price of the risky asset relative to the price $E_t m_{t,t+1}$ of a risk-free asset that pays off 1 for sure next period. The term $\left(\frac{\sigma_t(m_{t,t+1})}{E_t m_{t,t+1}}\right)$ is the 'market price of risk': it tells the rate at which the price ratio $a_t/E_t m_{t,t+1}$ deteriorates with increases in the conditional standard deviation of the pay out y_{t+1} .

Without imposing any theory about $m_{t,t+1}$, various studies have estimated the market price of risk $\left(\frac{\sigma_t(m_{t,t+1})}{E_t m_{t,t+1}}\right)$ from data on (a_t, y_{t+1}) For post WWII quarterly data, estimates of the market price of risk hover around .25. Hansen, XXX, and XXX's characterization of the equity premium puzzle is that .25 is much higher than would be implied by many theories that explicitly link $m_{t,t+1}$ to aggregate consumption. A standard benchmark is the theory $m_{t,t+1} = \beta u'(c_{t+1})/u'(c_t)$ where $u(\cdot)$ is a power utility function with power γ . That specification makes $m_{t,t+1} = \beta \left(\frac{c_{t+1}}{c_t}\right)^{\gamma}$. But aggregate consumption is a smooth series, so that the growth rate of consumption has a standard deviation so small that unless γ is implausibly large, the market price of risk implied by this theory of the stochastic discount factor $m_{t,t+1}$ remains far below the observed value of .25. Similarly, the permanent income model of chapter 9 that sets $m_{t,t+1} = M_c x_{t+1}/M_x x_t$ also implies too low a value of the market price of risk, again because the volatility of consumption growth is too small.

When the representative household is concerned about robustness, we have

$$a_t = E_t(m_{t,t+1}m_{t,t+1}^u)y_{t+1}$$
(12.5.3)

where from (12.4.9)

$$m_{t,t+1}^{u} = \exp\left[\epsilon_{t+1}' w_{t+1} - .5w_{t+1}' w_{t+1}\right].$$
 (12.5.4)

By construction, $E_t m_{t,t+1}^u = 1$. Hansen, Sargent, and Tallarini computed that $E_t (m_{t+1,t}^u)^2 = \exp(w_{t+1}'w_{t+1})$ so that

$$\sigma_t(m_{t,t+1}^u) = \sqrt{\exp(w_{t+1}'w_{t+1} - 1)} \approx |w_{t+1}'w_{t+1}|.$$
(12.5.5)

HST refer to $\sigma_t(m_{t,t+1}^u)$ as the one-period market price of Knightian uncertainty. Similarly, the $\tau - t$ -period market price of Knightian uncertainty is the conditional standard deviation of $m_{t,\tau}^u$ defined by (12.4.9). A fear of misspecification can boost the market price of risk by increasing these multiplicative adjustments to stochastic discount factors.



Figure 12.5.1: Market price of Knightian uncertainty for one-period securities $\sigma_t(m_{t,t+1})^u$ as function of detection error probability in HST model.



Figure 12.5.2: Market price of Knightian uncertainty for four-period securities $\sigma_t(m_{t,t+4})^u$ as function of detection error probability in HST model.

12.5.1. Calibrated market prices of Knightian uncertainty

As in chapter 9, we follow HST and use the parameterization $\sigma \equiv -\theta^{-1}$. HST computed one-period market prices of risk for a calibrated version of the permanent income model described in chapter 9. In particular, they proceeded as follows:

- 1. Setting $\sigma = 0$ and $\beta R = 1$, HST used the method of maximum likelihood to estimate the remaining free parameters of chapter 9's permanent income model.
- 2. HST used those maximum likelihood parameter estimates as the approximating model of the endowment processes d_t^*, \hat{d}_t for a representative agent whose continuation values they used to price risky assets. Thus, HST took a particular stand on how the representative agent created his approximating model, something that our robust control theory is silent about.
- 3. To study the effects of a fear of misspecification on asset prices while leaving the consumption-investment allocation (c_t, i_t) intact, HST lowered σ below zero, but adjusted the discount factor according to the relation $\beta = \hat{\beta}(\sigma)$ given by equation (9.3.18), which defines a locus of (β, σ) pairs that freeze $\{c_t, i_t\}$. For each (β, σ) thereby selected, HST calculated market prices of Knightian uncertainty and the detection error probabilities associated

with distinguishing the approximating model from the worst case model associated with σ . Figure 9.6.3 in chapter 9 reports those detection error probabilities as a function of σ . We are interested in the relation between the detection error probabilities and the *j*-period market prices of Knightian uncertainty.

4. For one and four period horizons, Figures 12.5.1 and 12.5.2 report the calculated market prices of Knightian uncertainty plotted against the detection error probabilities. These graphs have two salient features. First, there appear to be approximately linear relationships between the detection error probabilities and the market prices of Knightian uncertainty. In a continuous time, diffusion specification, Anderson, Hansen, and Sargent (2003) establish an exact linear relationship between the market price of risk and a bound on the detection error probabilities. To the extent that their bound is informative, their finding explains the striking pattern in these figures. Second, the market price of Knightian uncertainty is substantial even for values of the detection error probability sufficiently high that it seems plausible to seek robustness against models that close to the approximating model. Thus, a detection error probability of .3 leads to a one-period market price of uncertainty of about .15, which can explain about half of the observed equity premium.

In chapter 14, we shall return to the relationship between detection error probabilities and the market price of Knightian uncertainty in a version of a permanent income model in which the representative agent must use a Kalman filter because he does not observe the state variables that drive our two-factor endowment process.

12.6. A model of occupational choice and pay

Aloyisius Siow (1984) and Jaewoo Ryoo and Sherwin Rosen (2003) have used pure time-to-build structures to represent price and quantity cycles in markets for occupations under rational expectations. In their models, prospective new entrants into an occupation respond to optimal forecasts of the present value of a stream of wages that will begin accruing only after a period of schooling. We want to study how in equilibrium those forecasts and workers' decisions would behave under a concern for model misspecification.

Siow and Ryoo and Rosen used partial equilibrium models cast in terms of dynamic supply and demand curves. To analyze how a concern for model misspecification affects demand and supply, we first find the representative agent whose preferences induce the demand curve and the technology that generates the supply curve. It is straightforward to cast Ryoo and Rosen's model within the class of general equilibrium models of Chapter 10. Then the methods of section 12.2 and 12.4 can be used to construct a version of the model in which the representative agent has a concern about model misspecification indexed by $\theta \in (\underline{\theta}, \infty]$.

12.6.1. A one-occupation model

For concreteness, let the occupation be called engineering. Rosen (1995)'s model determines the stock of engineers N_t ; the number of new entrants into engineering school, n_t ; and the wage W_t of engineers. It takes k periods of schooling to become an engineer. We'll set k = 4 in our example. Ryoo and Rosen's model consists of the following equations: first, an inverse demand curve for engineers

$$W_t = \eta_d - \alpha_d N_t + u_{dt} , \ \alpha_d > 0; \tag{12.6.1}$$

second, a description of the education process as a time-to-build structure

$$N_{t+k} = \delta_N N_{t+k-1} + n_t , \ 0 < \delta_N < 1; \tag{12.6.2}$$

third, a definition of the expected present value of each new engineering student

$$v_t = \beta^k E_t \sum_{j=0}^{\infty} (\beta \delta_N)^j W_{t+k+j};$$
 (12.6.3)

and fourth, a supply curve of new students as a function of v_t

$$n_t = \eta_s + \alpha_s v_t + u_{st} , \ \alpha_s > 0. \tag{12.6.4}$$

Here $u_t = \begin{bmatrix} u_{dt} & u_{st} \end{bmatrix}'$ is a stochastic process of labor demand and supply shocks. Under a potentially distorted model indexed by w_{t+1} , the shocks u_{st}, u_{dt} are given by

$$u_{st} = U_s z_t$$

$$u_{dt} = U_{d1} z_t$$
(12.6.5)

where

$$z_{t+1} = A_{22}z_t + C_2(\epsilon_{t+1} + w_{t+1})$$
(12.6.6)

the eigenvalues of A_{22} are bounded in modulus by $\frac{1}{\sqrt{\beta}}$, U_s, U_{d1} are selector vectors, ϵ_{t+1} is an i.i.d. vector stochastic process with mean zero and covariance matrix I, and w_{t+1} is a vector of perturbations to the conditional means of the innovations to the approximating model. As usual, the approximating model assumes that $w_{t+1} \equiv 0$. Specification (12.6.5)–(12.6.6) allows the demand and supply shocks to be serially correlated.

We use the following:

Definition 12.6.1. A rational expectations equilibrium without a fear misspecification is a stochastic process $\{W_t, N_t, v_t, n_t\}_{t=0}^{\infty}$ satisfying (12.6.1), (12.6.2), (12.6.3), (12.6.4), (12.6.5) and (12.6.6), $w_{t+1} \equiv 0$, the stability condition $E_0 \sum_{t=0}^{\infty} \beta^t N_t^2 < +\infty$, and the initial conditions for $N_{-1}, n_{-s}, s = 1, \ldots, -k+1$.

12.6.2. Equilibrium with no concern about robustness

In the model in which the representative agent is not concerned about robustness, E_t in (12.6.3) is the mathematical expectation evaluated with respect to the distribution under the approximating ($w_{t+1} \equiv 0$) model. With a concern about robustness, the mathematical expectation under the distorted model \hat{E}_t replaces E_t in (12.6.3). But the distorted model is endogenous. Discovering it requires knowing the common preferences of the malevolent agent and the representative agent. To put a fear of misspecification into Ryoo and Rosen's model, it is thus necessary first to map the model without a fear of misspecification

306

into the equilibrium framework of chapter 10. This will identify the preferences of a representative agent that the malevolent agent also uses to formulate perturbations that promote robustness.

In terms of the class of general equilibrium models of chapter 10, we represent Ryoo and Rosen's model by sweeping the time-to-build structure into the household technology and the demand for engineers into the preference specification, while putting the supply of new engineers into the technology for producing goods. Here is how. Take the household technology to be

$$s_{t} = \begin{bmatrix} \alpha_{d} & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} N_{t} \\ n_{t-1} \\ n_{t-2} \\ n_{t-3} \end{bmatrix} + 0n_{t}$$

$$\begin{bmatrix} N_{t+1} \\ n_{t} \\ n_{t-1} \\ n_{t-2} \end{bmatrix} = \begin{bmatrix} \delta_{N} & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} N_{t} \\ n_{t-1} \\ n_{t-2} \\ n_{t-3} \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} n_{t}.$$

In the notation of chapter 10, these equations can be represented as 7

$$s_t = \Lambda h_{t-1} + \Pi c_t$$

$$h_t = \Delta_h h_{t-1} + \Theta_h c_t,$$
(12.6.7)

where we have set n_t in Ryoo and Rosen's model to c_t and $\begin{bmatrix} N_t & n_{t-1} & n_{t-2} & n_{t-3} \end{bmatrix}'$ to h_{t-1} in the model of chapter 10. To complete the representation of (12.6.1), we set the preference shock $b_t = \eta_d + u_{dt}$.

We represent the supply of entering students by using the technology side of the model. In particular, we assume

$$\begin{bmatrix} 1\\0 \end{bmatrix} c_t + \begin{bmatrix} -1\\\alpha_s^{-1} \end{bmatrix} i_t + \begin{bmatrix} 0\\-1 \end{bmatrix} g_t = \begin{bmatrix} 0\\0 \end{bmatrix} k_{t-1} + \begin{bmatrix} u_{st}\\0 \end{bmatrix}.$$

This equation matches the representation of technology in chapter 10

$$\Phi_c c_t + \Phi_i i_t + \Phi_g g_t = \Gamma k_{t-1} + d_t.$$

 $^{^{7}}$ In the language of Hansen and Sargent (200XX), this preference representation is not *canonical*, meaning that it must be transformed to a canonical representation in order to get convenient representations of dynamic demand functions.

Associated with this model is a representative agent who has preferences over c_t paths that are ordered by

$$-E_0 \sum_{t=0}^{\infty} \beta^t \left\{ .5(s_t - b_t) \cdot (s_t - b_t) + .5g_t \cdot g_t \right\}, \qquad (12.6.8)$$

where the mathematical expectation is taken with respect to the approximating model. Hansen and Sargent use the shadow prices from a planning problem to construct a competitive equilibrium, as described in chapter 10. The shadow prices $\mathcal{M}_t^s, \mathcal{M}_t^c, \mathcal{M}_t^h$ for s_t, c_t, h_t , respectively, satisfy

$$\mathcal{M}_t^s = b_t - s_t \tag{12.6.9a}$$

$$\mathcal{M}_t^h = E_t \sum_{\tau=1}^\infty \beta^\tau (\Delta_h')^{\tau-1} \Lambda' \mathcal{M}_t^s$$
(12.6.9b)

$$\mathcal{M}_t^c = \Theta_h' \mathcal{M}_t^h + \Pi' \mathcal{M}_t^s.$$
(12.6.9c)

Since $\Pi = 0$ for the present example, we have

$$\mathcal{M}_t^c = \Theta_h' E_t \sum_{\tau=1}^{\infty} \beta^\tau (\Delta_h')^{\tau-1} \Lambda' \mathcal{M}_t^s.$$
(12.6.10)

It can be verified that the wage W_t in Ryoo and Rosen's model matches the shadow price \mathcal{M}_t^s and that the present value v_t matches $\alpha_d^{-1}\mathcal{M}_{ct}$ (compare (12.6.3) with (12.6.10)). Where x_t is the state,⁸ Hansen and Sargent show that $\mathcal{M}_t^c = M_c x_t$ and $\mathcal{M}_t^s = M_s x_t$ and give formulas for the matrices M_c, M_s . We can use these objects to compute the equilibrium values of $W_t = M_s x_t, v_t =$ $M_c x_t$ in Ryoo and Rosen's model. The solutions for the quantities can be determined from the representation for the equilibrium in the state space form

$$x_{t+1} = A^o x_t + C\epsilon_{t+1}.$$

The next section computes examples of equilibria of the model both without and without a fear of misspecification. Appendix A solves the model by hand and describes some its analytical features.

⁸ The state x_t equals $[N_t \ n_{t-1} \ n_{t-2} \ 1 \ z_{st} \ z_{dt}]'$, where z_s is the supply shock and z_d is the demand shock. The presence of z_s and z_d means that we can accommodate demand and supply shocks that are first-order autoregressive processes.

12.6.3. Example

A version of the model with a fear of misspecification replaces (12.6.3) by

$$v_t = \beta^k \hat{E}_t \sum_{j=0}^{\infty} (\beta \delta_N)^j W_{t+k+j}, \qquad (12.6.11)$$

where \hat{E}_t is the mathematical expectation with respect to the distorted model. Representation (12.4.8) above implies that an equivalent representation of v_t in the model with a fear of misspecification is

$$v_t = \beta^k E_t \sum_{j=0}^{\infty} (\beta \delta_N)^j m_{t,t+k+j}^u W_{t+k+j}$$
(12.6.12)

where $m_{t,\tau}^u$ is the Radon-Nikodym derivative defined in (12.4.9). In (12.6.12), the expectation is evaluated under the approximating model. Equation (12.6.12) shows how a fear of misspecification puts an adjustment for model uncertainty into v_t . That adjustment gets reflected in the behavior of N_t, n_t, W_t in ways that the following example illustrates.

For alternative versions of the same model without a concern robustness (the solid lines) and with a concern for robustness with $-\theta^{-1} = -.5$ (the dotted lines), Fig. 12.6.1 shows impulse responses to an i.i.d. supply shock where the inverse demand shock is also i.i.d. Both of these impulse responses are evaluated under the approximating model.⁹ We set the covariance matrices of the two shocks to be I and the remaining parameter values at $\delta_N = .95$, $\alpha_s = 1$, $\alpha_d = .1$, $\eta_s = 10$, $\eta_d = 30$, $\beta = 1/1.05$.¹⁰

The effects of a fear of misspecification operate through the forecasting equation (12.6.11). The bottom left panel of Fig. 12.6.1 shows that when there is fear of misspecification, the initial adverse effect on v_t of a supply shock is greater in absolute value (more negative) than when there is no concern for robustness. The top right panel shows how, because of its more adverse implications for v_t , the supply shock causes a lower entry rates under a fear of

⁹ Thus, for the robust version of the model the agents inside the model are basing their decisions on the distorted model, but we are assuming that the data are actually generated by the approximating model.

¹⁰ See appendix A for a description of the role of the ratio $\alpha_s \alpha_d$ of the slopes inverse demand and supply function in influencing the solution, and for under the i.i.d. specification an inverse demand shock has no persistent effects on any variable in the model.

misspecification. This means that under the approximating model, the wage actually declines less in response to a supply shock under a fear of misspecification (see the top left panel). The top left panel shows wages declining less while the bottom left panel shows the expected present value declining more under a fear of misspecification. This discrepancy reflects the pessimistic forecasts that emanate from the worker's use of the distorted model to form \hat{E}_t . Wages decline less under a fear of misspecification because the lower entry rate induced by the pessimistic forecast v_t causes the actual stock of engineers N_t to increase less under a fear of misspecification (see the bottom right panel).¹¹



Figure 12.6.1: Impulse responses to supply shock without a fear of misspecification (solid lines) and with a fear of misspecification with $\sigma = -.5$.

 $^{^{11}}$ Appendix A gives analytical expressions that help provide more intuition about the shapes of the impulse response functions and the relations among them.

12.7. Concluding remarks

The asset pricing example of HST indicate how a little bit of concern about model misspecification can potentially substitute for a substantial amount of risk aversion when it comes to boosting theoretical values of market prices of risk. The boost in the market price of risk emerges from a form of pessimism relative to the representative agent's approximating model. The form that the pessimism takes is endogenous, depending both on the transition law and the representative agent's discount factor and one-period return function. Pessimism has been proposed by several researchers as an explanation of asset pricing puzzles, e.g., Reitz (XXXX) and Abel (20XXX). The contribution of the robustness framework is to discipline the appeal to pessimism by restricting the direction in which the approximating model is twisted, and by how much, through the detection probability statistics that we use to restrict θ .

A. Solving Ryoo and Rosen's model by hand

Using methods described by Sargent (1987), we can solve Ryoo and Rosen's model by hand and thereby discover a reduced description of the state. Substituting equations (12.6.1), (12.6.3), and (12.6.4) into (12.6.2) and rearranging yields

$$\left(1 + \beta \delta_N^2 + \alpha_s \alpha_d \beta^k - \delta_N L - \beta \delta_N L^{-1}\right) N_{t+k}$$

= $E_t \left[(1 - \beta \delta_N L^{-1})(\eta_s + u_{st}) + \alpha_s \beta^k (\eta_d + u_{d,t+k}) \right],$ (12.A.1)

where L is the backward shift operator. Notice the appearance in the characteristic polynomial of $\alpha_d \alpha_s = \frac{\alpha_d}{\alpha_s^{-1}}$, the ratio of the slope of the inverse demand schedule to the slope of the inverse supply schedule. The polynomial in L on the left side evidently can be factored as $f_0 f(\beta L^{-1}) f(L)$ where $f(L) = (1 - \psi L)$ and $|\psi| < 1$. Then the stabilizing solution of (12.A.1) is

$$N_{t+k} = \psi N_{t+k-1} + E_t \left\{ \left(\frac{f_0^{-1}}{1 - \psi \beta L^{-1}} \right) \left[(1 - \beta \delta_N L^{-1}) (\eta_s + u_{st}) + \alpha_s \beta^k (\eta_d + u_{d,t+k}) \right] \right\}$$
(12.A.2)

It follows from (12.A.2) that N_{t+k-1} is a complete description of the endogenous part of the state vector at the beginning of time t. We could have guessed this from (12.6.2) because N_{t+k-1} is independent of decisions or shocks that occur before time t.

When u_{st}, u_{dt} are i.i.d., (12.A.2) simplifies to

$$N_{t+k} = \psi N_{t+k-1} + \eta_s + \alpha_s \beta^k + u_{st}.$$
 (12.A.3)

In the i.i.d. case, it follows from (12.6.2) and (12.A.3) that the decision rule for n_t is

$$n_{t} = (\psi - \delta_{N})N_{t+k-1} + \eta_{s} + \alpha_{s}\beta^{k} + u_{st}.$$
 (12.A.4)
Part IV Robust filtering

Chapter 13 A robust filtering problem

13.1. Filtering

Chapter 4 described the Kalman filter, a recursive method for estimating a hidden state vector that can be computed as the dual of the optimal linear regulator. The Kalman filter assumes that a decision maker knows the statistical model linking the hidden state to observables. If he regards the statistical model only as approximating an unspecified true generating mechanism, he may want estimators of the hidden state that are robust to model misspecification. This chapter describes such robust estimators and how they embody another manifestation of Fellner's observation about how probability slanting depends on the 'prize' (see page 32): a robust filter depends partly on the decision maker's criterion function. In this chapter we assume that the decision maker cares about a weighted sum of current and *past* errors in estimating the state. An alternative but equivalent way of thinking about this criterion is that at some initial date, the decision maker commits himself to a particular estimator that he will eventually evaluate according to a weighted sum of state estimation errors over the entire horizon. Duality arguments that correspond to ones encountered in chapter 4 naturally lead us to this 'commitment' way of specifying the decision maker's criterion. Although this criterion is plausible for some economic problems, for others it is not, as we shall argue in chapter 14. We adopt it now partly because it of how it illuminates the duality of robust control and filtering.

– 315 –

13.1.1. Warning about recycling of notation

As is typical in some presentations of filtering and control, we have recycled some notation. For example, we use the matrices A and K to denote different objects in the robust filtering and control problems.

13.2. Summary of robust filtering and duality

The duality described in chapter 4 doubles the usefulness of the optimal linear regulator problem because to each control problem there corresponds a filtering problem, and vice versa. There is also a filtering problem that is dual to the robust linear regulator problem studied in chapters 6 and 7. This section displays that problem but postpones a formal derivation of it to section 13.3. We begin by recalling the duality of control and filtering presented in chapter 4, then indicate how it can be extended to incorporate concerns about model misspecification.

13.2.1. Duality of ordinary filtering and control

Let $\tilde{z}_t = C'\lambda_t + D'\mu_t$. Consider the optimal linear regulator problem

$$-\lambda_0' \Sigma \lambda_0 = \max_{\{\mu_t\}} \sum_{t=0}^{\infty} -\tilde{z}_t' \tilde{z}_t$$
(13.2.1)

where the maximization is subject to an initial condition for λ_0 and the transition law

$$\lambda_{t+1} = A'\lambda_t + G'\mu_t. \tag{13.2.2}$$

Here λ_t is the state vector and μ_t is the control vector.

Chapter 4 displayed a filtering problem that is dual to this regulator problem and interpreted the λ_t 's and μ_t 's of the control problem as Lagrange multipliers associated with that filtering problem. For $t \ge 0$, the filtering problem has a state vector x_{-t} and an observation vector y_{-t} that satisfy

$$x_{-t} = Ax_{-t-1} + C\epsilon_{-t} \tag{13.2.3a}$$

$$y_{-t} = Gx_{-t-1} + D\epsilon_{-t} \tag{13.2.3b}$$

where ϵ_{-t} is an i.i.d. Gaussian vector with mean zero and covariance matrix I. The recursive estimator of the hidden state is:

$$\hat{x}_{-t} = A\hat{x}_{-t-1} + K(y_{-t} - G\hat{x}_{-t-1}), \qquad (13.2.4)$$

where K is the Kalman gain matrix. The error in reconstructing the state at t is

$$e_{-t} = x_{-t} - \hat{x}_{-t}. \tag{13.2.5}$$

The decision maker wants to construct an estimator \hat{z}_{-t} of a linear combination $z_{-t} \equiv Hx_{-t}$ of the state at each t and so poses the minimization problem

$$\min_{K} E \lim_{T \to \infty} T^{-1} \sum_{t=0}^{T} (z_{-t} - \hat{z}_{-t})' (z_{-t} - \hat{z}_{-t})$$
(13.2.6)

or

$\min_{K} \operatorname{trace}(H'H\Sigma)$

subject to (13.2.3), (13.2.4), (13.2.5) and where again $\Sigma = E(e_{-t}e'_{-t})$. The Σ that minimizes trace $(H'H\Sigma)$ is independent of H. Define the operators

$$\mathcal{K}(\Sigma) = (CD' + A\Sigma G')(DD' + G\Sigma G')^{-1}$$
(13.2.7)

$$T^{*}(\Sigma) = (A - \mathcal{K}(\Sigma)G)\Sigma(A - \mathcal{K}(\Sigma)G)'$$

+ $(G - \mathcal{K}(\Sigma)D)(G - \mathcal{K}(\Sigma)D)'$ (12.2.8)

$$+ (C - \mathcal{K}(\Sigma)D)(C - \mathcal{K}(\Sigma)D)'.$$
(13.2.8)

The minimized value of Σ solves the Riccati equation

$$\Sigma = T^*(\Sigma) \tag{13.2.9}$$

and the minimizing K satisfies

$$K = \mathcal{K}(\Sigma). \tag{13.2.10}$$

In chapter 4, we established that (K, Σ) also solve the linear regulator (13.2.1), (13.2.2), and that $\mu_t = -K'\lambda_t$ is the optimal decision rule.

13.2.2. A robust linear regulator

We now describe a multiplier robust control problem that corresponds to (13.2.1), (13.2.2) and whose dual is a robust filtering problem. For $\theta \in (\underline{\theta}, +\infty]$, consider an undiscounted robust linear regulator corresponding to (13.2.1), (13.2.2):

$$-\lambda_0' \Sigma \lambda_0 = \max_{\{\mu_t\}} \min_{\{\phi_{t+1}\}} \sum_{t=0}^{\infty} \left\{ -\tilde{z}_t' \tilde{z}_t + \theta \phi_{t+1}' \phi_{t+1} \right\}$$
(13.2.11)

where again $\tilde{z}_t = C'\lambda_t + D'\mu_t$ and the maximization is subject to

$$\lambda_{t+1} = A'\lambda_t + G'\mu_t + H'\phi_{t+1}$$
(13.2.12)

and an initial condition λ_0 . The solution of this robust control problem is a pair of decision rules

$$\mu_t = -K'\lambda_t \tag{13.2.13a}$$

$$\phi_{t+1} = K'_{\phi} \lambda_t. \tag{13.2.13b}$$

13.2.3. A dual robust filtering problem

We can now state the filtering problem that is dual to (13.2.11), (13.2.12). The robust filtering problem surrounds the approximating model (13.2.3) with a set of perturbed models of the form

$$x_{-t} = Ax_{-t-1} + C(\epsilon_{-t} + w_{-t}) \tag{13.2.14a}$$

$$y_{-t} = Gx_{-t-1} + D(\epsilon_{-t} + w_{-t}) \tag{13.2.14b}$$

where ϵ_{-t} is another i.i.d. Gaussian vector with mean zero and covariance matrix I and w_{-t} is a vector of measurable functions of $[y^{-t}, x^{-t}]$. The w_{-t} process represents specification errors that feed back on histories of both the unobserved state and the observed variables, thereby representing misspecified dynamics. The decision maker constructs a robust filter by solving the following two-player zero-sum multiplier game:¹

$$\operatorname{trace}(H'H\Sigma) = \max_{\{w_{-t}\}_{t=0}^{\infty}} \min_{K} E \lim_{T \to \infty} T^{-1} \sum_{t=0}^{T} \left(z'_{-t} z_{-t} - \theta w'_{-t} w_{-t} \right) \quad (13.2.15)$$

¹ When H is such that the decision maker cares about unobserved components of the state, the solution to this problem will depend on H, so that the Σ associated with the solution depends on H (which it did not for the ordinary Kalman filter).

where $z_{-t} = He_{-t}$ and subject to (13.2.14), (13.2.4), (13.2.5) and where $\Sigma = E(e_{-t}e'_{-t})$. For H'H of full rank, the extremizing Σ is the fixed point of iterations on $T^* \circ \mathcal{D}^*$, where T^* is the operator (13.2.10) associated with iterations on the Riccati equation for the ordinary Kalman filter, and

$$\mathcal{D}^*(\Sigma) = \Sigma + \theta^{-1} \Sigma H' (I - \theta^{-1} H \Sigma H')^{-1} H \Sigma.$$
(13.2.16)

Lars: is there a *direct* (in terms of filtering) interpretation of this as a distorted expectations operator? The only interpretation I could figure out was a dual one. The robust Kalman gain K then satisfies

$$K = \mathcal{K} \circ \mathcal{D}^*(\Sigma), \tag{13.2.17}$$

where

$$\Sigma = T^* \circ \mathcal{D}^*(\Sigma) \tag{13.2.18}$$

and \mathcal{K} is defined in (13.2.7). The maximizing w_{-t} sequence has the recursive representation

$$w_{-t} = -\theta^{-1} \left[I + \theta^{-1} (C - KD)' \Sigma^{-1} (C - KD) \right]^{-1} (C - KD)' \Sigma^{-1} (A - KG) e_{-t-1},$$

which shows how the worst case mean distortions w feed back on $e_{-t-1} = x_{-t-1} - \hat{x}_{-t-1}$, which means that they feed back on both the unobserved state x_{-t-1} and the history of the observed signal as summarized by \hat{x}_{-t-1} . This formula shows how the malevolent (i.e., error-*maximizing*) agent exploits his information advantage over the decision maker, who cannot observe x_{-t-1} .

13.2.4. The robust filter when Hx_{-t} is observable

Formulas (13.2.16), (13.2.17), and (13.2.18) show that the robust filter is context specific, depending on the criterion $z'_{-t}z_{-t} = x'_{-t}H'Hx_{-t}$ that the maximizing player cares about. In many problems, $x_{-t}H'Hx_{-t}$ plays the role of the maximizing player's one-period return function. When the decision maker cares only about observable components of the state, formulas (13.2.16), (13.2.17), and (13.2.18) imply that $\mathcal{D}^*(\Sigma) \equiv \Sigma$ because $H\Sigma H' = 0$. In this case, the robust filter is equivalent with the ordinary Kalman filter.

13.3. The robust filtering problem

We now substantiate these claims about the robust filter and its duality with the robust linear regulator. Because chapters 6 and 7 study discounted optimal linear regulators, we extend (13.2.11), (13.2.12) to make it into a discounted optimal linear regulator, and seek a filtering problem that is dual to that discounted problem.

As we saw in chapter 4, though the filtering problem is typically applied in stochastic contexts, because the mathematics merely manipulates moment matrices, we can present what seems to be a nonstochastic derivation of a robust Kalman filter. Thus, we shall appeal to a version of certainty equivalence to drop the random process ϵ_{-t} from the state space system and also shall drop mathematical expectations from the prediction error criterion that the decision maker seeks to minimize. We let the random vector $\epsilon_{-t} \equiv 0$ in the decision maker's approximating model and allow for specification errors to be of the form:²

$$x_{-t} = Ax_{-t-1} + Cw_{-t} \tag{13.3.1a}$$

$$y_{-t} = Gx_{-t-1} + Dw_{-t}.$$
(13.3.1b)

Here $t \ge 0$. Let y^{-t} denote the history of y up to -t. Let $\hat{E}[\cdot|y^{-t-1}]$ denote a filtered value conditioned on the history of y up to time -t - 1. We seek filtered values $\hat{x}_{-t} = \hat{E}[x_{-t}|y^{-t}], \ \hat{y}_{-t} = \hat{E}[y_{-t}|y^{-t-1}]$, where $\hat{E}(\cdot)$ is a distorted expectations operator. Note the differences in the conditioning sets. See (13.3.2) and (13.3.4) for confirmation of these different conditioning sets.

We construct a robust filter by solving a non-stochastic zero-sum two-player game in which an evil prediction-error-maximizing agent chooses a sequence of shocks $\{w_{-t}\}$ to maximize the same prediction error criterion that a decision maker wants to minimize. We restrict the robust Kalman gain to be timeinvariant. For now, we take as given the choice K by the prediction-errorminimizing decision maker, and focus on the decision of an evil maximizing agent. To set the problem facing the evil agent, we form an observer system (see Kwakernaak and Sivan (1972)). Emulating the measurement equation, we require the estimator of y to take the form:

$$\hat{y}_{-t} = G\hat{x}_{-t-1}.\tag{13.3.2}$$

 $^{^2}$ See the discussion in section 4.2 for an explanation for why we make time recede into the past with increases in t.

It follows that the prediction error for y_{-t} is:

$$y_{-t} - \hat{y}_{-t} = G(x_{-t-1} - \hat{x}_{-t-1}) + Dw_{-t}.$$

We consider updating schemes for \hat{x} that are parameterized by a fixed gain matrix K. The forecast-error-minimizing agent chooses K in an updating rule of the form:³

$$\hat{x}_{-t} = A\hat{x}_{-t-1} + K(y_{-t} - \hat{y}_{-t})$$
(13.3.3)

or

$$\hat{x}_{-t} = A\hat{x}_{-t-1} + K(y_{-t} - G\hat{x}_{-t-1}).$$
(13.3.4)

Subtracting (13.3.3) from (13.3.1a) gives

$$x_{-t} - \hat{x}_{-t} = (A - KG)(x_{-t-1} - \hat{x}_{-t-1}) + (C - KD)w_{-t}.$$
 (13.3.5)

Define the state reconstruction error:

$$e_{-t} = x_{-t} - \hat{x}_{-t}. \tag{13.3.6}$$

Then (13.3.5) can be expressed

$$e_{-t} = (A - KG)e_{-t-1} + (C - KD)w_{-t}.$$
(13.3.7)

The filter K is designed to minimize a quadratic form in the following linear combination of the forecast errors in the state:

$$z_{-t} = He_{-t}.$$
 (13.3.8)

³ The reader of Başar and Bernhard (1995) will notice that their predictor (e.g, 6.62 on page 273) appears to have a somewhat different form than ours, there being an extra term in \hat{x}_t in theirs. However, if we apply their formulas to the problem in our text, it can be shown that that extra term vanishes. The relevant part of their analysis occurs on their pages 272 and 273. Their approach is to solve two Riccati equations. Take the first one, that in M_k on page 272. Define $G_k = -I$ which makes the target Hx - u so that we have a pure forecasting problem where the goal is to forecast Hx, and $B = 0, \bar{A} = A$, and $\bar{Q} = 0$, so that the Riccati equation for M implies that $M_k = 0$ as a solution when the equation is stated in a form that does not require nonsingularity. The full information control problem has a zero value function and the optimal control is to set u = Hx, so that the control portion of the problem is degenerate. The nontrivial part of their solution in our case is the Γ_k component and Başar and Bernhard refer to this as the dual quantity. Their equation for Σ_{kk} will have a nondegenerate solution. Substituting $\hat{u}_k = H\check{x}$ into their equations (6.62) and (6.63) gives the result that the extra \check{x} term drops out of (6.62).

The criterion of the multiplier form of a two-person game for a robust filter is

$$.5\sum_{t=0}^{\infty}\beta^{t}(z_{-t}'z_{-t} - \theta w_{-t}'w_{-t}), \quad \beta \in (0,1)$$
(13.3.9)

Notice how β^t weights forecast errors from the more recent past more heavily. The decision maker minimizes this criterion by choosing K, while the evil agent maximizes it by choosing w_{-t} 's subject to (13.3.7) and (13.3.8). Here $\theta > 0$ penalizes the $w'_{-t}w_{-t}$ sequence. Given K, the maximized value of (13.3.9) is

$$.5e_0'\Sigma^{-1}e_0 \tag{13.3.10}$$

where Σ satisfies a Riccati equation for a dual problem to be stated shortly and the maximizing w_{-t} sequence can be represented in the recursive form $w_{-t} = \mathcal{O}e_{-t-1}$, where e_{-t} evolves according to (13.3.7) and where a formula for \mathcal{O} will be given in (13.3.13) below.

13.3.1. The evil agent's problem

We now focus on the evil agent's problem. As in chapter 7, we can use the optimal value function that emerges from this problem as a criterion function that the minimizing agent can use to devise a robust K. Recall that in posing the problem of the evil agent in chapter 7, we frequently took the control law u = -Fx as fixed and let the evil agent respond. We now take K as fixed and study the problem of maximizing (13.3.9) by choice of $\{w_t\}$. We form the conjugate problem associated with choosing the w_{-t} 's to maximize (13.3.9). Let $\beta^t \lambda'_t$ denote the vector of Lagrange multipliers on (13.3.7), let $\beta^t \phi'_t$ be the vector of multipliers on (13.3.8), and form a Lagrangian. Among the first-order conditions for the problem of maximizing the Lagrangian with respect to $\{w_{-t}, e_{-t}\}_{t=0}^{\infty}$ and minimizing it with respect to $\{\lambda_t, \phi_t\}_{t=0}^{\infty}$ are:

$$w_{-t}: \quad w_{-t} = -\frac{1}{\theta} (C - KD)' \lambda_t$$
 (13.3.11*a*)

$$z_{-t}: \quad z_{-t} = -\phi_t \tag{13.3.11b}$$

$$e_{-t-1}: \quad \beta \lambda_{t+1} = (A - KG)' \lambda_t + \beta H' \phi_{t+1}$$
 (13.3.11c)

$$e_{-0}: \quad \lambda_0 = H'\phi_0 \tag{13.3.11d}$$

We can use (13.3.11a) to get a convenient formula for the distortion w_{-t} . First, note that λ_t is the vector of shadow prices of e_{-t} , so that $\lambda_t = \Sigma^{-1}e_{-t}$ where Σ appears in the value function (13.3.10). This equation and (13.3.11a) imply

$$w_{-t} = -\theta^{-1} (C - KD)' \Sigma^{-1} e_{-t}.$$
 (13.3.12)

To compute the feedback rule for the worst-case shock w_{-t} , substitute (13.3.7) into (13.3.12) and solve for w_{-t} to get

$$w_{-t} = -\theta^{-1} \left[I + \theta^{-1} (C - KD)' \Sigma^{-1} (C - KD) \right]^{-1} (C - KD)' \Sigma^{-1} (A - KG) e_{-t-1}$$
(13.3.13)

Below, we shall show how both Σ and K can be computed by solving the linear regulator (13.2.11), (13.2.12).

13.3.2. The dual to the evil agent's problem

To form the dual to the error-maximizing agent's problem, use (13.3.11a) and (13.3.11b) to write

$$w'_{-t}w_{-t} = \frac{1}{\theta^2}\lambda'_t(C - KD)(C - KD)'\lambda_t$$
$$z'_{-t}z_{-t} = \phi'_t\phi_t.$$

Substituting these into (13.3.9) gives the dual criterion

$$\frac{1}{2\theta} \sum_{t=0}^{\infty} \beta^t \left\{ -\lambda_t' (C - KD) (C - KD)' \lambda_t + \theta \phi_t' \phi_t \right\}.$$
(13.3.14)

The dual problem is to minimize (13.3.14) by choice of $\{\phi_t\}_{t=0}^{\infty}$, subject to (13.3.11*c*) and (13.3.11*d*). For convenience, rewrite (13.3.11*c*), (13.3.11*d*) as

$$\lambda_{t+1} = \beta^{-1} (A - KG)' \lambda_t + H' \phi_{t+1}$$
(13.3.15a)

$$\lambda_0 = H'\phi_0. \tag{13.3.15b}$$

This is a discounted linear regulator problem with state λ_t and control ϕ_{t+1} . The optimized value of the objective functions of the original and dual problems are equal. We can reinterpret the dual problem in terms of a time-domain version of the multiplier problem associated with (7.2.3), (7.2.4), which for convenience we repeat here:

$$\min_{\{w_t\}} \sum_{t=0}^{\infty} \beta^t \left\{ -x'_t (H_0 - JF)' (H_0 - JF) x_t + \theta w'_t w_t \right\}$$
(13.3.16)

subject to

$$x_{t+1} = (A_o - BF)x_t + Cw_{t+1}$$
(13.3.17*a*)

$$x_0 = Cw_0. (13.3.17b)$$

Notice that the dual filtering problem (13.3.14), (13.3.15a), (13.3.15b) corresponds to (13.3.16), (13.3.16), (13.3.17) with the settings in Table 13.3.1. This means that all of the computational methods that apply to the control problem can be used to solve the filtering problem, as we describe in the following section.

Table 13.3.1: Duality of filter and control

Filter	Σ	ϕ_t	λ_t	A'/eta	G'/β	H'	C'	K'	D'
Control	P	w_t	x_t	A_o	В	C	H_0	F	J

13.3.3. Computing K

By using the duality relations listed in Table 13.3.1, we can formulate a twoperson zero-sum game that can be used to compute a robust filter K. We simply use analogies from chapters 6 and 7 to take advantage of what we know about the analogous game for the control problem.

The two-player zero-sum game associated with the robust control problem is

$$\max_{\{u_t\}} \min_{\{w_t\}} \sum_{t=0}^{\infty} \beta^t \left\{ -z'_t z_t + \theta w'_t w_t \right\}$$
(13.3.18)

subject to

$$z_t = H_0 x_t + J u_t \tag{13.3.19a}$$

$$x_{t+1} = A_o x_t + B u_t + C w_{t+1} (13.3.19b)$$

$$x_0 = Cw_0. (13.3.19c)$$

The u_t component of the solution is a time-invariant feedback rule

$$u_t = -Fx_t, \tag{13.3.20}$$

where formulas for F are given in chapters 2, 6, and 7.

Using Table 13.3.1, it follows that the two-player zero-sum game for the dual filtering problem is

$$\max_{\{\mu_t\}} \min_{\{\phi_t\}} \sum_{t=0}^{\infty} \beta^t \{ -\tilde{z}'_t \tilde{z}_t + \theta \phi'_t \phi_t \}$$
(13.3.21)

subject to

$$\tilde{z}_t = C'\lambda_t + D'\mu_t \tag{13.3.22a}$$

$$\lambda_{t+1} = \beta^{-1} A' \lambda_t + \beta^{-1} G' \mu_t + H' \phi_{t+1}$$
(13.3.22b)

$$\lambda_0 = H'\phi_0. \tag{13.3.22c}$$

This problem (13.3.21), (13.3.22) can also be formulated as an optimal linear regulator. Equilibrium choices of μ_t, ϕ_{t+1} have representations of the forms given in (13.2.13), namely, $\mu_t = -K'\lambda_t$ and $\phi_{t+1} = K'_{\phi}\lambda_t$, where K', K'_{ϕ} can be calculated with formulas analogous to those used to solve the corresponding control problem (13.3.18), (13.3.19).

The beauty of these findings is that we can solve for the robust Kalman gain by solving a robust optimal linear regulator problem.

13.3.4. Matlab programs

We can induce our matlab program doublex9.m for solving the robust linear regulator to compute K, Σ , and $\mathcal{D}^*(\Sigma)$ by exploiting the duality displayed in section 13.3.3. We accomplish this in the following steps.⁴

 Input the objects A, C, G, D that form the state-space system, H and β in the decision maker's criterion function, and θ; A is n×n, C is n×p, G is m×n, D is m×p, and H is r×n, where n is the dimension of the state, p the number of shocks, m the number of observables, and r the number of variables entering the decision maker's criterion function.

 $^{^4\,}$ The Matlab program <code>rfilter.m</code> performs these steps. .

- 2. Prepare the robust linear regulator formulated in (13.3.21) and (13.3.22) by setting $Q = CC', R = DD', W = CD', A = \beta^{-1}A', B = \beta^{-1}G', D = H'$. These are the objects in a discounted robust linear regulator problem with cross-products between states and controls. Therefore, if we want to use doublex9.m, which is for undiscounted problem without cross-products between states and controls, we have to use the trick described in chapter 3 on page 65 for converting to such a problem. We accomplish this in step 3.
- 3. Form $As = \sqrt{\beta}(A BR^{-1}W')$, $Bs = \sqrt{\beta}B$, $Qs = Q WR^{-1}W'$. The Matlab program trick.m accomplishes these tasks.
- 4. Set sig = $-\theta^{-1}$ and issue the command [F,K,P,Pt]=doublex9(As,Bs,D,Qs,R,sig). To complete the trick begun in step 3, set $\bar{F} = F + R^{-1}W$, as described in chapter 3 on page 66.
- 5. Finally, $K = \overline{F}$, $\Sigma = P$, and $\mathcal{D}^*(\Sigma) = Pt$.

Given K, the feedback rule for the worst case mean w can be computed by using either formula (13.3.13) or (13.A.2). Appendix A describes an optimal linear regulator problem that can be used to compute this feedback rule and that leads to formula (13.A.2). Appendix B describes a worst-case distorted transition law for the state for which the ordinary Kalman gain equals the robust Kalman filter for the transition law under the approximating model.

13.4. Robustifying Muth's filter

As an example, we set $\beta = 1$ and consider Muth's (1960) problem of estimating the position of a random walk disturbed by measurement error. We set H = 1(so that the decision maker cares about the hidden state), and assume the approximating model:

$$x_{t+1} = x_t + \alpha \hat{\epsilon}_{1,t+1} \tag{13.4.1a}$$

$$y_{t+1} = x_t + \hat{\epsilon}_{2,t+1} \tag{13.4.1b}$$

where α is the signal to noise ratio and $\hat{\epsilon}_{t+1} = \begin{bmatrix} \hat{\epsilon}_{1,t+1} & \hat{\epsilon}_{2,t+1} \end{bmatrix}'$ is an i.i.d. Gaussian process with mean zero and identity covariance matrix. The state x_t

326

is to be estimated from current and past values of y_t . Setting H = 1 makes the decision maker's criterion equal to the variance of the error in reconstructing the state x from past signals y. We consider the filter

$$\hat{x}_{t+1} = \hat{x}_t + K(y_{t+1} - \hat{x}_t) \tag{13.4.2}$$

where \hat{x}_{t+1} is the estimate of the state using the history of y_s through t+1. We want K to be robust to misspecification of (13.4.1).

We use the program rfilter.m described in section 13.3.4 to compute the robust K. To illuminate the perturbations to the approximating model that the decision maker is considering, we also display the following calculations.

To attain robustness, the decision maker considers a family of perturbed models:

$$x_{t+1} = x_t + \alpha(\epsilon_{1,t+1} + w_{1,t+1}) \tag{13.4.3a}$$

$$y_{t+1} = x_t + \epsilon_{2,t+1} + w_{2,t+1} \tag{13.4.3b}$$

where $\epsilon_{t+1} = \begin{bmatrix} \epsilon_{1,t+1} & \epsilon_{2,t+1} \end{bmatrix}'$ is another i.i.d. Gaussian process with mean zero and identity covariance matrix, and $\begin{bmatrix} w_{1,t+1}, w_{2,t+1} \end{bmatrix}$ are distortions to the conditional means of the two shocks $\hat{\epsilon}_{t+1}$ in (13.4.1). Subtracting (13.4.2) from (13.4.3*a*) and using (13.4.3*b*) gives

$$e_{t+1} = (1 - K)e_t + \alpha \epsilon_{1,t+1} - K\epsilon_{2,t+1} + \alpha w_{1,t+1} - Kw_{2,t+1}, \qquad (13.4.4)$$

where $e_t \equiv x_t - \hat{x}_t$. Using formula (13.3.13), we can represent the worst case mean distortions as

$$w_{1,t+1} = -N_1 e_t$$

$$w_{2,t+1} = -N_2 e_t,$$
(13.4.5)

where N_1 and N_2 are computed using formula (13.3.13) or (13.4.2). Please notice that N_1, N_2 are functions of θ and K.

For arbitrary K and fixed $w_{1,t+1} = -N_1e_t, w_{2,t+1} = -N_2e_t$, the error in reconstructing the state when the model associated with (N_1, N_2) prevails is

$$e_{t+1} = (1-K)e_t - \alpha N_1 e_t + K N_2 e_t + \alpha \epsilon_{1,t+1} - K \epsilon_{2,t+1}$$
(13.4.6)

or

$$e_{t+1} = \chi e_t + \alpha \epsilon_{1,t+1} - K \epsilon_{2,t+1}, \qquad (13.4.7)$$

where

$$\chi = 1 - K - \alpha N_1 + K N_2. \tag{13.4.8}$$

Equation (13.4.7) gives the law of motion of the error e_t in reconstructing the state for filter K when the conditional means of the shocks are feeding back on e_t via N_1, N_2 . Denote the variance of e_t by $\operatorname{var}_e(K; N_1, N_2)$. From (13.4.7) it follows directly that

$$\operatorname{var}_{e}(K; N_{1}, N_{2}) = \frac{\alpha^{2} + K^{2}}{1 - \chi^{2}}.$$
 (13.4.9)

The spectral density of e is

$$S_e(\omega; K, N_1, N_2) = g_1(\omega)g_1(-\omega) + g_2(\omega)g_2(-\omega)$$
(13.4.10)

where $g_1(\omega) = \frac{\alpha}{1-\chi \exp(-i\omega)}$, $g_2(\omega) = \frac{K}{1-\chi \exp(-i\omega)}$; S_e achieves the decomposition of var_e across frequencies:

$$\operatorname{var}_e = \frac{1}{2\pi} \int_{-\pi}^{\pi} S_e(\omega; K, N_1, N_2) d\omega.$$

Consider var_e as a function of K. Let $\hat{K}(\theta)$ be the robust filter associated with θ . When $N_1(\theta), N_2(\theta)$ deliver the worst case distortions w_1 and w_2 to the conditional means of the two components of ϵ for a given θ , $\operatorname{var}_e(K; N_1, N_2)$ is minimized at $K = \hat{K}(\theta)$.

13.4.1. Ordinary Kalman filter

Let $K^* = \hat{K}(+\infty)$ denote the standard Kalman filter. If $\theta = +\infty$, then $N_1 = N_2 = 0$ and the variance of e_t simplifies to:

$$\operatorname{var}_{e}(K;0,0) = \frac{\alpha^{2} + K^{2}}{1 - (1 - K)^{2}}$$

$$= \frac{\alpha^{2} + K^{2}}{2K - K^{2}}.$$
(13.4.11)

Minimizing (13.4.11) with respect to K gives a formula for K that agrees with that produced by the ordinary Kalman filter: $K^* = \frac{\sqrt{\alpha^4 + 4\alpha^2 - \alpha^2}}{2}$.⁵

328

⁵ When $\alpha = 1$, this equals $\frac{\sqrt{5}-1}{2}$, the golden ratio.



Figure 13.4.1: Variance of $e_t(K; N_1, N_2)$ as function of K for N_1 and N_2 evaluated at $\theta = 10^8$. Here the ordinary Kalman gain K^* satisfies $K^* \approx \hat{K}(\theta)$, and both K^* and $\hat{K}(\theta)$ are denoted by asterisks. The two curves are for two values of the signal-noise ratio $\alpha = 1$ and $\alpha = 1.78$.



Figure 13.4.2: Variance of $e_t(K; N_1(\theta), N_2(\theta))$ as function of K for N_1 and N_2 evaluated at $\theta = 7$. Here the ordinary Kalman gain K^* satisfies $K^* < \hat{K}(\theta)$ (where \hat{K} is denoted by the x and K^* by the small vertical line on the $\operatorname{var}_e(K)$ curves). The two curves are for two values of the signal-noise ratio $\alpha = 1, 1.78$.

13.4.2. Illustrations



Figure 13.4.3: Frequency decomposition of the reconstruction error variance $\operatorname{var}_e(K; N_1(\theta), N_2(\theta))$ for $\theta = 10^8$ for $\hat{K}(\theta)$ and K^* , $\alpha = 1$. The two curves for $\hat{K}(\theta)$ and K^* approximately coincide.



Figure 13.4.4: Frequency decomposition of the reconstruction error variance $\operatorname{var}_e(K; N_1(\theta), N_2(\theta))$ for $\theta = 7$ for $\hat{K}(\theta)$ and K^* , $\alpha = 1$. The solid curve is for the robust gain \hat{K} , the dotted one for the ordinary Kalman gain K^* . The robust gain \hat{K} flattens the decomposition of variance across frequencies.



Figure 13.4.5: Frequency decomposition of the reconstruction error variance $\operatorname{var}_e(K; N_1, N_2)$ for $\theta = 7$ for $\hat{K}(\theta)$ and K^* , $\alpha = 1.78$. The solid curve is for the robust gain \hat{K} , the dotted one for the ordinary Kalman gain K^* .



Figure 13.4.6: The robust Kalman gain $\hat{K}(\theta)$ as a function of $\log(\theta)$ and α .

For fixed $\theta < \infty$, we can determine $\hat{K}(\theta)$ by solving the two player game (13.3.21), (13.3.22). We can also find the associated feedback rules for the shocks $N_1(\theta), N_2(\theta)$ using formula (13.3.13).

In Fig. 13.4.1, we have fixed $\theta = 10^8$ and derived the associated \hat{K}, N_1, N_2 (all three are functions of θ) and have plotted $\operatorname{var}_e(K; N_1, N_2)$, the variance of $e_t(K)$, as a function of K. It has a minimum at $\hat{K}(\theta)$. We have also



Figure 13.4.7: The robust Kalman gain $\hat{K}(\theta)$ as a function of $\log(\theta)$, given $\alpha = 1.78$.

put $K^* = \hat{K}(+\infty)$ and $\hat{K}(\theta)$ on the graph. For this large value of θ , K^* is indistinguishable from $\hat{K}(\theta)$.

Fig. 13.4.2 contains the same information as Fig. 13.4.1, except for the value of $\theta = 7$. Now $\hat{K}(7) > K^* = K(\infty)$, though the state reconstruction error variances var_e associated with them are close.

Fig. 13.4.3 displays the frequency decomposition of $\operatorname{var}_e(K^*; 0, 0)$. Because $N_1 = N_1 = 0$, this is the frequency decomposition of the variance of e_t under the assumption of no specification error, using the ordinary Kalman gain K^* with $\alpha = 1$. (See (13.4.7), (13.4.8).) Fig. 13.4.4 displays the frequency decomposition of $\operatorname{var}_e(K; N_1(7), N_2(7))$ for two values of K: K^* and $\hat{K}(7)$. Here 7 is the value of θ . Thus, the dotted line is the frequency decomposition of $\operatorname{var}_e(\hat{K}^*; N_1(7), N_2(7))$, while the solid line is the frequency decomposition of $\operatorname{var}_e(\hat{K}(7); N_1(7), N_2(7))$. Because they are computed using (13.4.7), (13.4.8) evaluated at $N_1(7), N_2(7)$, these spectral densities describe the frequency decompositions of the variances of the reconstruction errors associated with K^* and $\hat{K}(\theta)$ under the worst case model associated with $\theta = 7$.⁶ Fig. 13.4.4 is for $\alpha = 1$, while Fig. 13.4.5 if for $\alpha = 1.78$. Note that for the ordinary Kalman gain K^* , the spectral density under the approximating model in Fig. 13.4.3 is lower at low frequencies than is the spectral density in Fig. 13.4.4, which is evaluated under the worst case model. This illustrates how the evil agent spends

⁶ Tom: check to make sure the interpretation of the Ball and permanent income models are correct w.r.t. which model is used to evaluate the spectrum.

most of his 'entropy budget' by deceiving the decision maker at low frequencies. Fig. 13.4.4 shows how the robust filter responds by lowering the low frequency contributions to variance under the worst case model.

Fig. 13.4.3 shows that the ordinary Kalman filter K^* is most vulnerable to low frequency components of e_t , which can be induced by having the worst case conditional means feed back positively on e_t . Fig. 13.4.4 shows how the worst case conditional means associated with $\theta = 7$ pump up the low frequencies of e_t , and how the robust $\hat{K}(7)$ filter achieves a lower variance $\operatorname{var}_e(K; N_1, N_2)$ by flattening the spectrum, accepting higher variance at higher frequencies in exchange for lower variance at the low frequencies where the worst case conditional means operate the strongest.

Fig. 13.4.6 and Fig. 13.4.7 show the robust Kalman gain \hat{K} as functions of $\log(\theta)$ and α . These figures show how increasing the preference for a robust filter (i.e., decreasing θ) raises the Kalman gain.

13.5. Another example

Velde (2005XXXX) considers a state-space system of the form (13.3.1) in which the first component of x_{-t} is core-inflation and remaining components describe the dynamics of a vector of relative prices. Velde is interested in minimizing $E(x_{1,-t}-\hat{x}_{1,-t})^2$. To get a robust Kalman filter, we would take Velde's estimates of A, C, G, D and set $H = \begin{bmatrix} 1 & 0 & \cdots & 0 \end{bmatrix}$.

13.6. Discounting and the direction of time

This chapter has been partly motivated by mechanical questions associated with duality. The duality of ordinary (non-robust) filtering and control described in chapter 4 led us to suspect that there exists a filtering problem that is dual to the robust discounted optimal linear regulator problem analyzed in chapters 6 and 7. By reverse engineering, we have found that robust filtering problem.

This robust filter minimizes a geometrically discounted sum of current and *past* forecast errors. That criterion is inappropriate in economic models that tell the decision maker to care only about current and *future* returns, settings that

induce the decision maker to limit his attention to errors in forecasting current and *future* variables that influence payoffs.

In chapter 14, we shall describe a different robust filtering problem with a payoff function that is exclusively forward looking. Nevertheless, it is possible to imagine some economic contexts in which the timing convention of the filter of this chapter can be defended, for example, where the agent who filters must commit himself in advance to a filtering rule and then be judged in terms of the average forecasting behavior only after much time has elapsed. The formulation studied in this chapter is widely used in the control literature (for example, see Başar and Bernhard (1995) and Whittle (1990)).

13.7. Another perturbation

Consider the situation in which the decision maker cares only about current and future values of the state. In particular, suppose that, conditional on knowing the state x he has a value function -x'Mx, where the symmetric and positive semi-definite matrix M might be obtained as the outcome of iterating on a Riccati equation that looks to the future. Suppose that if the decision maker knew the state, he would use the decision rule u = -Fx. Now suppose that the state is not known and that the decision maker estimates the state using an ordinary Kalman filter. The Kalman filter makes the time t state x distributed according to a Gaussian distribution with conditional mean \hat{x} and conditional covariance Σ , so that the decision maker acts as if the (partially hidden) state x obeys $x = \hat{x} + \hat{e}$, where \hat{e} is normal with mean 0 and covariance Σ . An application of a certainty equivalence argument would have the decision maker use the decision rule $u = -F\hat{x}$.

But now suppose that the decision maker believes that the posterior distribution emerging from the Kalman filter is misspecified, so that x instead obeys $x = \hat{x} + \hat{e} + u$ where u is a perturbation to the conditional mean. Because the decision maker wants decisions that are robust with respect to such misspecifications, he conducts the following context-specific worst case analysis, where now the 'context' inspires him to choose the distortion u to harm the forward-looking criterion -x'Mx. To find a worst case perturbation u, penalize u by

entropy as measured by $u\Sigma^{-1}u$ and consider the problem

$$\min_{u} - (\hat{x} + u)' M(\hat{x} + u) + \frac{\theta}{2} u \Sigma^{-1} u$$
(13.7.1)

whose first order necessary condition implies

$$u = -(M - \theta \Sigma^{-1})^{-1} M \hat{x}.$$
 (13.7.2)

Apply the partitioned inverse formula

$$(a - bd^{-1}c)^{-1} = a^{-1} + a^{-1}b[d - ca^{-1}b]^{-1}ca^{-1}$$

with $a = M, b = \theta, d = \Sigma, c = 1$ to get

$$(M - \theta \Sigma^{-1})^{-1} = M^{-1} + M^{-1} \theta [\Sigma - M^{-1} \theta]^{-1} M^{-1}.$$

Then from (13.7.2) we have

$$u = -\left(I + M^{-1}[\theta^{-1}\Sigma - M^{-1}]^{-1}\right)\hat{x}.$$

The worst case x is then $\check{x} = \hat{x} + u$ or

$$\check{x} = (I - \theta^{-1} \Sigma M)^{-1} \hat{x}.$$
(13.7.3)

The decision maker achieves robustness to doubts about the specification of the prior distribution coming from the Kalman filter by using the decision rule

$$u = -F\check{x}.\tag{13.7.4}$$

13.7.0.1. Formulation from control theory literature

Başar and Bernhard (1995) and Whittle (1990) use a decision rule of the form (13.7.4) where instead of starting with the prior for x that emerges from the Kalman filter, they begin from the distorted prior that we deduced from our reverse-engineered-from-duality robust Kalman filter. In that formulation, *past* distortions w_{-t} to the conditional means of past ϵ_{-t} 's, used to design the robust Kalman filter, affect the decision rule through their effects on \hat{x} ; while *future* distortions w_t to conditional means of future ϵ_t 's, affect the decision rule through the design of a robust F via the approach described in chapter 2 and 6. See Whittle (1990) and Başar and Bernhard (1995) for extensive discussions.

13.7.1. Transition to next chapter

The material in this section forms a bridge to the next chapter, where we consider a 'two- θ recursiv formulation of decision problems in which the decision makers considers two sources of possible misspecification: (1) misspecified dynamics, accommodated by allowing distortions to the conditional mean of the state and measurement errors ϵ_{t+1} to feed back on past outcomes, and a misspecified prior distribution for the hidden state variables, accommodated by a perturbation like the *u* in problem (13.7.1). By using different θ 's to penalize the entropies associated with these two perturbations, we can focus the decision maker's concerns more on one or the other of these sources of misspecification.

A. Another way to compute the worst-case shock

Formula (13.3.13) shows how to compute the worst-case shock associated with the robust filter. An alternative way to compute it is to use the dual problem to compute K, then to formulate the primal problem, say with the following convention for our time index t:

$$\max_{\{w_t\}} \sum_{t=0}^{\infty} \beta^{-t} (e_t' H' H e_t - \theta w_t' w_t)$$
(13.A.1a)

subject to

$$e_t = (A - KG)e_{t-1} + (C - KD)w_t$$
(13.A.1b)

given an initial e_{-1} . Note how the discounting of the past in problem (13.3.9), (13.3.7), (13.3.8) corresponds to anti-discounting the future in (13.A.1a) (because

 $|\beta| < 1$). The feedback rule for the worst case shock is

$$w_t = Ne_{t-1},$$
 (13.A.2)

where N equals the usual feedback matrix 7 for the optimal linear regulator associated with (13.A.1).

Having found N either from formula (13.3.13) or (13.A.2), we return to our original convention for the index t and use (13.3.2) to represent the robust filter recursively as

$$\hat{y}_{-t} = G\hat{x}_{-t-1} \tag{13.A.3a}$$

$$\hat{x}_{-t} = A\hat{x}_{-t-1} + K(y_{-t} - \hat{y}_{-t}).$$
(13.A.3b)

The associated worst-case law of motion for the state and the observed variables is

$$w_{-t} = N(x_{-t-1} - \hat{x}_{-t-1}) \tag{13.A.4a}$$

$$x_{-t} = Ax_{-t-1} + Cw_{-t} \tag{13.A.4b}$$

$$y_{-t} = Gx_{-t-1} + Dw_{-t}.$$
 (13.A.4c)

B. A Bayesian interpretation of the robust *K*

In section 6.5.2 of chapter 6, we remarked that when a decision maker solves an ordinary optimal linear regulator problem, acting as if a particular representation of the worst case law of motion is true, he will thereby compute the robust decision rule. By displaying a particular law of motion for the state that renders the robust rule optimal, that finding was the basis for a Bayesian interpretation of the robust control law. In this appendix, we describe a corresponding Bayesian interpretation of the robust filter by displaying a representation of the worst-case law of motion for which an ordinary Kalman gain will be the robust Kalman gain. We accomplish this by using another version of a "Big X, little x trick.

Let $K = K(\theta)$ be the robust Kalman gain and express the law of motion for $\{x_t, y_t\}$ under the worst case model as

$$w_{t+1} = N(x_t - \check{x}_t)$$
(13.B.1a)
$$x_{t+1} = Ax_t + C(\epsilon_{t+1} + w_{t+1})$$
(13.B.1b)

$$\begin{aligned} x_{t+1} &= Ax_t + C(\epsilon_{t+1} + w_{t+1}) \\ \check{r}_{t+1} &= A\check{r}_t + K(u_{t+1} - C\check{r}_t) \end{aligned} \tag{13.B.10}$$

$$x_{t+1} = Ax_t + K(y_{t+1} - Gx_t)$$
(13.B.1c)
$$y_{t+1} = Gx_t + D(e_{t+1} + w_{t+1})$$
(13.B.1d)

$$y_{t+1} = Gx_t + D(\epsilon_{t+1} + w_{t+1}).$$
(13.B.1d)

⁷ Namely, -F in $u_t = -Fx_t$ in the linear regulator in chapters 2 and 6.

We use $(\check{\cdot})$ to replace $(\hat{\cdot})$ for a reason that we explain after presenting system (13.B.2). Substituting (13.B.1a) into (13.B.1b) and (13.B.1d), substituting the resulting formula for y_{t+1} into (13.B.1c), and rearranging allows us to represent the worst case law as

$$\begin{bmatrix} x_{t+1} \\ \check{x}_{t+1} \end{bmatrix} = \begin{bmatrix} A+CN & -CN \\ -K(G+DN) & A+K(G+DN) \end{bmatrix} \begin{bmatrix} x_t \\ \check{x}_t \end{bmatrix} + \begin{bmatrix} C \\ KD \end{bmatrix} \epsilon_{t+1}$$

$$y_{t+1} = \begin{bmatrix} (G+DN) & -DN \end{bmatrix} \begin{bmatrix} x_t \\ \check{x}_t \end{bmatrix} + D\epsilon_{t+1}.$$
(13.B.2)

Notice how this representation shows that the worst-case law of motion is no longer Markov in the state x_t alone, though it is Markov when the state is augmented to include \check{x}_t . This occurs because the distorted mean w_{t+1} feeds back on $x_t - \check{x}_t$, thereby inducing history-dependence into the law of motion for x_t . Express system (13.B.2) as

$$X_{t+1} = \bar{A}X_t + \bar{C}\epsilon_{t+1} y_{t+1} = \bar{G}X_t + D\epsilon_{t+1}.$$
(13.B.3)

Apply the ordinary Kalman filter to this system to get the Kalman gain \check{K} and the representation \hat{X} ~

$$\ddot{X}_{t+1} = \bar{A}\ddot{X}_t + \dot{K}(y_{t+1} - \bar{G}\ddot{X}_t).$$
 (13.B.4)

Then $\check{K} = K(\theta)$, where $K(\theta)$ is the robust Kalman filter.

Chapter 14 Joint control and estimation

In commerce bygones are forever bygones and we are always starting clear at each moment, judging the value of things with a view to future utility. Industry is essentially prospective not retrospective.

— William Stanley Jevons, 1871

14.1. Introduction

This chapter extends ideas about control and filtering from chapters 4, 6, and 13. We study a decision maker who does not observe parts of the state that help forecast variables that he cares about. We formulate a joint control and prediction problem and show how it can be represented recursively.

The filtering problem in chapter 13 took as the decision maker's approximating model a state space representation for observables and states, posed the problem of estimating a function of hidden states with a filter that is robust to perturbations of the approximating model. The approach in this chapter has a different starting point. We take the view that the decision maker's approximating model *includes* a recursive representation of the estimator of the hidden state that is derived by applying the ordinary (nonrobust) Kalman filter to the approximating state space model for states and observables. This means that we include the sufficient statistics for the distribution of the hidden part state that come from the ordinary Kalman, namely the mean and covariance of the hidden part of the state, as among the state variables. We then obtain robust decision rules by allowing an evil agent to perturb conditional distributions for that approximating model. Among other distributions, the evil agent perturbs the distribution of the hidden state conditional on the history of signals. The decision maker designs a decision rule that is robust to these perturbations.

This chapter draws heavily from Hansen and Sargent (2005b XXXX).

- 339 -

14.2. A recursive control and filtering problem

14.2.1. The decision maker's approximating model

Following Hansen and Sargent (2005b XXXX), Partition a state vector as

$$x_t = \begin{bmatrix} y_t \\ z_t \end{bmatrix}$$

where y_t is observed and z_t is not observed by a decision maker whose oneperiod utility function is

$$U(x_t, a_t) = -.5 \begin{bmatrix} a'_t & x'_t \end{bmatrix} \begin{bmatrix} Q & P \\ P' & R \end{bmatrix} \begin{bmatrix} a_t \\ x_t \end{bmatrix}$$

where a_t is a vector of controls that influence future values of both (y, z) and a signal s that is informative about z The decision maker orders $\{x_t, a_t\}$ sequences according to

$$E\left[\sum_{t=0}^{\infty} \beta^t U(x_t, a_t) \middle| y_0\right], \qquad (14.2.1)$$

where E is the mathematical expectation with respect to a probability distribution that we now describe. At time t + 1, the decision maker observes a vector s_{t+1} that includes y_{t+1} and possibly other signals about the hidden state. The decision maker remembers past signals. The variables obey

$$y_{t+1} = \Pi_s s_{t+1} + \Pi_y y_t + \Pi_a a_t \tag{14.2.2a}$$

$$z_{t+1} = A_{21}y_t + A_{22}z_t + B_2a_t + C_2w_{t+1}$$
(14.2.2b)

$$s_{t+1} = D_1 y_t + D_2 z_t + H a_t + G w_{t+1} \tag{14.2.2c}$$

where $w_{t+1} \sim \mathcal{N}(0, I)$. Substituting (14.2.2*c*) into (14.2.2*a*) gives the following transition law for the observed state:¹

$$y_{t+1} = A_{11}y_t + A_{12}z_t + B_1a_t + C_1w_{t+1}, (14.2.3)$$

¹ Sometimes we formulate a problem directly in terms of (14.2.3) without first stipulating (14.2.2a).

where $A_{11} = (\Pi_s D_1 + \Pi_y), A_{12} = \Pi_s D_2, B_1 = (\Pi_s H + \Pi_a), C_1 = \Pi_s G$. Thus, we have the state-space system

$$x_{t+1} = Ax_t + Ba_t + Cw_{t+1} \tag{14.2.4a}$$

$$s_{t+1} = Dx_t + Ha_t + Gw_{t+1}.$$
 (14.2.4b)

The decision maker believes that the distribution of the initial value of the unobserved part of the state is

$$z_0 \sim \mathcal{N}(\check{z}_0, \Delta_0). \tag{14.2.5}$$

Let $s^t = [s_t, \ldots, s_0]$. Taking into account that y_t is observed and applying the ordinary Kalman filter to system (14.2.4) gives the following representation for $\tilde{z}_{t+1} \equiv E[z_{t+1}|s^{t+1}]$ and $\tilde{s}_{t+1} \equiv E[s_{t+1}|s^t]$:

$$\check{s}_{t+1} = D_1 y_t + D_2 \check{z}_t + H a_t \tag{14.2.6a}$$

$$\check{z}_{t+1} = A_{21}y_t + A_{22}\check{z}_t + B_2a_t + K_2(\Delta_t)(s_{t+1} - \check{s}_{t+1})$$
(14.2.6b)

$$\Delta_{t+1} = A_{22}\Delta_t A'_{22} + C_2 C'_2 - K_2(\Delta_t) (A_{22}\Delta_t D'_2 + C_2 G')' \quad (14.2.6c)$$

$$K_2(\Delta) = (A_{22}\Delta D'_2 + C_2 G')(D_2 \Delta D'_2 + G G')^{-1}$$
(14.2.6d)

where $\Delta_t = E[z_t - \check{z}_t][z_t - \check{z}_t]'$. Notice that \check{z}_{t+1} conditions on s^{t+1} and that \check{s}_{t+1} conditions on s^t . Under the approximating model, $z_t \sim \mathcal{N}(\check{z}_t, \Delta_t)$. Thus, (\check{z}_t, Δ_t) are sufficient statistics for the distribution $q_t(z_t)$ of the unobserved part of the state at t.

We regard representation (14.2.6) as a complete statement of the decision maker's approximating model. Thus, we take the laws of motion for (\check{z}_t, Δ_t) that come from applying the Kalman filter to model (14.2.4) to be parts of the decision maker's approximating model.² We shall design a robust decision rule by considering the robustness of a decision rule to statistical perturbations of (14.2.6). This structure identifies distributions to two random vectors to perturb at date t: (1) the conditional distribution of the shock w_{t+1} , which according to the approximating model is $\mathcal{N}(0, I)$; and (2) the distribution of the hidden state z_t , which according to the approximating model is $\mathcal{N}(\check{z}_t, \Delta_t)$.³

 $^{^2}$ This contrasts with the approach in chapter 13, where the approximating model did not include the law of motion for an estimate of the hidden state induced by applying the ordinary Kalman filter.

 $^{^{3}}$ A virtue of this formulation is that by taking the outcomes of the Kalman filter as applied to (14.2.4) as part of the approximating model, we have to solve only one Kalman filtering

14.2.2. Two sources of statistical perturbation

We transform representation (14.2.6) in a way designed to focus our attention on perturbations to the distributions of w_{t+1} and z, respectively. To formulate a recursive version of our problem, let * denote a next period value and use (14.2.6) to express the evolution equation for $[y'_{t+1} \quad \check{z}'_{t+1} \quad \Delta_{t+1}]$ as

$$\begin{bmatrix} y^* \\ \check{z}^* \\ \Delta^* \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} & 0 \\ A_{21} & A_{22} & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} y \\ \check{z} \\ f(\Delta) \end{bmatrix} + \begin{bmatrix} B_1 \\ B_2 \\ 0 \end{bmatrix} a + \begin{bmatrix} A_{12} \\ K_2(\Delta)D_2 \\ 0 \end{bmatrix} [z-\check{z}] + \begin{bmatrix} C_1 \\ K_2(\Delta)G \\ 0 \end{bmatrix} w^{\prime}$$
(14.2.7)

where $f(\Delta) = A_{22}\Delta A'_{22} - (A_{22}\Delta D'_2 + C_2G')(D_2\Delta D'_2 + GC')^{-1}(A_{22}\Delta A'_{22} + G_2C')' + C_2C'_2$, $w^* \sim \mathcal{N}(0, I)$, and $z \sim \mathcal{N}(\check{z}, \Delta)$. Notice that Δ_t evolves exogenously with respect to (y, \check{z}) , so that given an initial condition Δ_0 , a path $\{\Delta_{t+1}\}_{t=0}^{\infty}$ can be computed before observing any data.

Two random vectors appear in representation (14.2.7): $(z - \tilde{z})$ and w^* . In the next subsection, we describe systematic ways of organizing context-specific perturbations of the distributions of these two random vectors. At first reading, it is possible lightly to skim these subsections and move immediately to subsection 14.2.5 where we use a certainty equivalent problem to find the key objects needed to compute a robust decision rule.

14.2.3. Two operators

An operator \mathbf{T}^1 systematically perturbs the distribution of w_{t+1} conditional on (y, \check{z}, z) and another operator \mathbf{T}^2 perturbs the distribution of z conditional on (y, \check{z}) . Let $q = (\check{z}, \Delta)$ denote our sufficient statistics for the distribution of the hidden state. Throughout this section, we let a be a measurable function of (y, q).

problem. Alternative formulations can be conceived that would have the decision maker solve a separate filtering problem for each perturbation of the approximating model (14.2.4). Those would obviously be less convenient computationally. See Hansen and Sargent (2005a, 2005b XXX) for a discussion of this issue.

14.2.3.1. The T^1 operator

Let m be a nonnegative random variable that is measurable with respect to (y^*, q^*, z^*) . Assume that m has mean 1 conditional on (y, q, z, a). Hansen and Sargent (2005a, 2005b XXXX) show that m can be used to represent distortions to the joint distribution of (y^*, q^*, z^*) , conditional on (y, q, z, a). In particular, m functions as a Radon-Nikodym derivative, or likelihood ratio, for transforming one distribution into another. Let V(y, q, z, a) be a measurable function of (y, q, z). Then $E\left[mV(y^*, q^*, z^*, a)|y, z, q, a\right]$ equals the conditional expectation of V evaluated with respect to a distorted density formed by multiplying the original density by m. Define the entropy of m by $\varepsilon^1(m) = E(m \log m | y, q, z)$. Following Hansen and Sargent (2005bXXXX), we define the operator \mathbf{T}^1 by:

$$(\mathbf{T}^{1}V(y^{*}, q^{*}, z^{*}, a)|\theta)(y, q, z, a) = \min_{m \ge 0} E\Big[mV(y^{*}, q^{*}, z^{*}, a)|y, z, q, a\Big] + \theta\varepsilon^{1}(m)$$
(14.2.8a)

$$= \theta \log \int \exp\left[\frac{-V(y^*, q^*, z^*, a)}{\theta}\right] \phi(w^*) dw^*, \qquad (14.2.8b)$$

where ϕ is the standard normal density and the minimization in (14.2.8*a*) is subject to the law of motion (14.2.7) and the restriction that E[m|y, z, q, a] = 1. The minimizing *m* is

$$m^{\heartsuit} \propto \exp\left[\frac{-V(y^*, q^*, z^*, a)}{\theta}\right],$$
 (14.2.9)

where the factor of proportionality is chosen to make the mean of m conditional on (y, q, a) equal to 1.

14.2.3.2. The T^2 operator

Let h be a nonnegative random variable that is a function of (y,q,z,a) and that has mean 1 conditional on (y,q,a). Hansen and Sargent (2005a, 2005b XXXX) show that such a nonnegative random variable h can be used to represent distortions to the joint distribution of (y,q,z). Define the entropy of has $\varepsilon^2(h) = E(h \log h | y, q, a)$. Consider a function $\check{V}(y,q,z,a)$ and define the operator

$$(\mathbf{T}^{2}\check{V}(y,q,z,a)|\theta)(y,q,a) = \min_{h\geq 0} E\Big(h\check{V}(y,q,z,a)|y,q,a\Big) + \theta\varepsilon^{2}(h)$$
(14.2.10a)

$$= -\theta \log \int \exp\left[\frac{-\check{V}(y,q,z,a)}{\theta}\right] q(z)dz \qquad (14.2.10b)$$

where the minimization is subject to E[h|y,q,a] = 1. The minimizing h is

$$h^{\heartsuit} \propto \exp\left[\frac{-\dot{V}(y,q,z,a)}{\theta}\right],$$
 (14.2.11)

where the factor of proportion is chosen to make h have expectation equal to 1.

14.2.4. A recursive formulation for control and estimation

By finding a worst-case m, the operator \mathbf{T}^1 distorts the distribution of w^* , conditional on (y, q, z, a). It can help a decision maker explore fragility of decisions to misspecification of the distribution of (y^*, q^*, z^*) conditional on (y, q, z, a). Notice that the hidden state z is included in the conditioning set.

By finding a worst-case h, the operator \mathbf{T}^2 distorts the distribution of z conditional on (y, q, a). It can help a decision maker explore the fragility of a decision rule $a = a(y, \tilde{z})$ to misspecifications of the prior distribution for z, as represented in our linear quadratic problem, for example, by the sufficient statistics (\tilde{z}, Δ) for the distribution q(z).

By solving the following Bellman equation, the decision maker can design a decision rule that is robust to misspecifications of the conditional distribution of w^* and the distribution of the hidden state z:

$$W(y,q) = \max_{a} \mathbf{T}^{2} \Big[U(x,a) + \mathbf{T}^{1} [\beta W^{*}(y,q) | \theta_{1}](y,q,z,a) | \theta_{2}] \Big] (y,q,a).$$
(14.2.12)

Here \mathbf{T}^1 integrates over w^* , conditioning on (y, q, z, a), and \mathbf{T}^2 integrates over z, conditioning on (y, q, z). Assigning different values to θ in the two operators lets the decision maker focus more on misspecifications of one rather than the other of the two distributions being perturbed.

By virtue of equations (14.2.8a) and (14.2.10a), each step in recursion (14.2.12) involves maximization over a and minimization over h and m. For linear quadratic problems like the ones we are interested in here, certainty equivalence ideas can be exploited to simplify the calculations. We show how to do this in the next section by focusing first on the means of the perturbed distributions of w^* and z, then calculating the covariance matrices later.

14.2.5. A certainty equivalent short cut

Let u be the mean of $z - \check{z}$ and \tilde{v} the mean of w^* , both conditioned on (y, \check{z}) . Consider the evolution equation

$$\begin{bmatrix} y^* \\ \tilde{z}^* \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} y \\ \tilde{z} \end{bmatrix} + \begin{bmatrix} B_1 \\ B_2 \end{bmatrix} a + \begin{bmatrix} A_{12} \\ K_2(\Delta)D_2 \end{bmatrix} u + \begin{bmatrix} C_1 \\ K_2(\Delta)G \end{bmatrix} \tilde{v} \quad (14.2.13)$$
or
$$\tilde{z}^* = \tilde{A}\tilde{z} + \tilde{P}(\Delta)\tilde{z} \quad (14.2.14)$$

$$\check{x}^* = \tilde{A}\tilde{x} + \tilde{B}(\Delta)\tilde{a} \tag{14.2.14}$$

where

$$\tilde{a} = \begin{bmatrix} a \\ u \\ \tilde{v} \end{bmatrix}, \qquad \tilde{x} = \begin{bmatrix} y \\ \check{z} \end{bmatrix}.$$
(14.2.15)

We obtained (14.2.13) by letting $w^* = \tilde{v} + \epsilon^*$ and $z - \check{z} = u + \epsilon_u$, where ϵ^* and ϵ_u are both Gaussian random vectors with means of zero, and then dropping the terms in ϵ^* and ϵ_u from (14.2.7). Dropping these terms can be justified by a certainty equivalence argument like one used in chapter 2. We will take account of these omitted terms later when we compute covariance matrices. The important thing to note now is that omitting these terms at this stage does not imperil our computations of u and \tilde{v} .

We add to the utility function U(x, a) the parts of two entropy terms pertinent for our deterministic problem to get the augmented return function

$$U(x,a) + \theta_1 |\tilde{v}|^2 + \theta_2 u' \Delta^{-1} u_2$$

where θ_1 penalizes distortions v and θ_2 penalizes distortions u. Define this augmented return function as $r(\tilde{x}, \tilde{a})$ and represent it as

$$\begin{split} r(\tilde{x}, \tilde{a}) &= -\frac{1}{2} \begin{bmatrix} a' & y' & z' \end{bmatrix} \begin{bmatrix} Q & P_1 & P_2 \\ P_1' & R_{11} & R_{12} \\ P_2' & R_{21} & R_{22} \end{bmatrix} \begin{bmatrix} a \\ y \\ z \end{bmatrix} + \frac{\theta_1}{2} |\tilde{v}|^2 + \frac{\theta_2}{2} u' \Delta^{-1} u \\ &= -\frac{1}{2} \begin{bmatrix} \tilde{a}' & \tilde{x}' \end{bmatrix} \Pi(\Delta) \begin{bmatrix} \tilde{a} \\ \tilde{x} \end{bmatrix} \end{split}$$

and

$$\Pi(\Delta) = \begin{bmatrix} \Pi_{11} & P_{12} \\ \Pi_{21} & \Pi_{22} \end{bmatrix}$$

where

$$\Pi_{11} = \begin{bmatrix} Q & P_2 & 0 \\ P'_2 & R_{22} - \theta_2 \Delta^{-1} & 0 \\ 0 & 0 & -\theta_1 I \end{bmatrix}, \quad \Pi_{12} = \begin{bmatrix} P_1 & P_2 \\ R_{21} & R_{22} \\ 0 & 0 \end{bmatrix}, \quad \Pi_{22} = \begin{bmatrix} R_{11} & R_{12} \\ R_{21} & R_{22} \end{bmatrix}$$

and $\Pi_{21} = \Pi'_{12}$. Then we can compute (y, \check{z}) -contingent distortions to the means (u, \tilde{v}) and a robust decision rule for *a* that solves (14.2.12) by solving the deterministic problem:

$$\min_{\{u_t, \tilde{v}_t\}} \max_{\{a_t\}} \sum_{t=0}^{\infty} \beta^t r(\tilde{x}_t, \tilde{a}_t)$$
(14.2.16)

subject to

$$\tilde{x}_{t+1} = \tilde{A}\tilde{x}_t + \tilde{B}(\Delta_t)\tilde{a}_t \tag{14.2.17}$$

where Δ_t is the solution of (14.2.6*c*), (14.2.6*d*). The Bellman equation for problem (14.2.16), (14.2.17) is

$$-\frac{1}{2}\tilde{x}'\Omega(\Delta)\tilde{x} = \operatorname{ext}_{\tilde{a}}\left\{-\frac{1}{2}\begin{bmatrix}\tilde{a}' & \tilde{x}'\end{bmatrix}\Pi(\Delta)\begin{bmatrix}\tilde{a}\\\tilde{x}\end{bmatrix} - \beta\tilde{x}'\Omega(\Delta^*)\tilde{x}^*\right\},\qquad(14.2.18)$$

where ext denotes extremization (i.e., maximization with respect to a and minimization with respect to u and \tilde{v}) and the extremization is subject to (14.2.14). To be well posed, (θ_1, θ_2) must be large enough to make the matrix

$$\begin{bmatrix} \theta_{2}\Delta^{-1} - R_{22} & 0\\ 0 & \theta_{1}I \end{bmatrix} - \beta \begin{bmatrix} A_{12}' & D_{2}'K_{2}(\Delta)'\\ C_{1}' & G'K_{2}(\Delta)' \end{bmatrix} \Omega(\Delta^{*}) \begin{bmatrix} A_{12} & C_{1}\\ K_{2}(\Delta)D_{2} & K_{2}(\Delta)G \end{bmatrix}$$
(14.2.19)

be positive definite. We call (14.2.19) a 'no breakdown condition'.

The decision rule for \tilde{a} is

$$\tilde{a} = -\left[\Pi_{11}(\Delta) + \beta \tilde{B}(\Delta)' \Omega^*(\Delta^*) \tilde{B}(\Delta)\right]^{-1} \left[\Pi_{12} + \beta \tilde{B}(\Delta)' \Omega^*(\Delta^*) \tilde{A}\right] \tilde{x}$$
(14.2.20)

and the recursion in $\Omega(\Delta)$ is the Riccati equation

$$\Omega(\Delta) = \Pi_{22} + \beta \tilde{A}(\Delta)' \Omega^*(\Delta^*) \tilde{A}(\Delta) - \left[\Pi_{12} + \beta \tilde{B}(\Delta)' \Omega^*(\Delta^*) \tilde{A}\right]' \\ \left[\Pi_{11}(\Delta) + \beta \tilde{B}(\Delta)' \Omega^*(\Delta^*) \tilde{B}(\Delta)\right]^{-1} \left[\Pi_{12} + \beta \tilde{B}(\Delta)' \Omega^*(\Delta^*) \tilde{A}\right].$$
(14.2.21)

In the special case in which the decision maker in effect conditions on an infinite history of signals and in which Δ_t has converged, we can set $\Delta^* = \Delta$ and exploit the observation that, as noted in chapter 2, problem (14.2.18) can be solved using standard formulas for the ordinary discounted optimal linear regulator. In particular, our Matlab program olrp.m can be applied.

The extremizing decision rule from either (14.2.20) in the general case or olrp.m for the special case $\Delta_{t+1} = \Delta_0 \forall t \ge 1$ is

$$\tilde{a} \equiv \begin{bmatrix} a \\ u \\ \tilde{v} \end{bmatrix} = -\tilde{F}(\Delta)\tilde{x} = -\begin{bmatrix} \tilde{F}_1(\Delta) \\ F_2(\Delta) \\ F_3(\Delta) \end{bmatrix} \tilde{x}$$
(14.2.22)

The first block row gives the robust decision rule and the second gives the distorted mean u of $z - \tilde{z}$ both as functions of $\tilde{x} = \begin{bmatrix} y \\ \tilde{z} \end{bmatrix}$. The third block row gives the mean \tilde{v} of the distorted distribution for w^* , conditional on \tilde{x} .

In problem (14.2.12), the distorted mean actually depends on the unobserved state z as well, since the \mathbf{T}^1 operator conditions on z. So while problem (14.2.16), (14.2.17) allows us to compute the decision rule a and the distortion to the mean of z that solves (14.2.12), it does not provide the mean distortion to the distribution of w^* conditioned on (y, \tilde{z}, z) that the \mathbf{T}^1 operator in (14.2.12) computes. We now describe how to compute the distorted mean of w^* conditioned on the set (y, \tilde{z}, z) that conditions \mathbf{T}^1 . We want this conditional mean and the associated conditional covariance matrix in order to compute other objects that will be of interest later.

14.2.6. Computing the T^1 operator

Impose the robust control law for a in the law of motion (14.2.14) to get

$$\begin{bmatrix} y^*\\ \tilde{z}^* \end{bmatrix} = \tilde{A} \begin{bmatrix} y\\ \tilde{z} \end{bmatrix} - \begin{bmatrix} B_1\\ B_2 \end{bmatrix} \tilde{F}_1(\Delta) \begin{bmatrix} y\\ \tilde{z} \end{bmatrix} + \begin{bmatrix} A_{12}\\ K_2(\Delta)D_2 \end{bmatrix} [z-\tilde{z}] + \begin{bmatrix} C\\ K_2(\Delta)G \end{bmatrix} w^*$$

or
$$\begin{bmatrix} y^*\\ \tilde{z}^* \end{bmatrix} = \bar{A}(\Delta) \begin{bmatrix} y\\ \tilde{z} \end{bmatrix} \bar{H}(\Delta)[z-\tilde{z}] + \bar{C}(\Delta)w^*.$$

We already know the mean of the worst case w^* conditioned on (y, \tilde{z}) . We want to know the mean of w^* conditioned on (y, \tilde{z}, z) . An LQ control problem associated with the \mathbf{T}^1 operator is:

$$\min_{v} -\frac{1}{2}\beta \begin{bmatrix} y^{*\prime} & \check{z}^{*\prime} \end{bmatrix} \Omega^{*}(\Delta^{*}) \begin{bmatrix} y^{*} \\ \check{z}^{*} \end{bmatrix} + \frac{\theta_{1}}{2}v^{\prime}v$$

where the minimization is subject to the law of motion

$$\begin{bmatrix} y^*\\ \check{z}^* \end{bmatrix} = \bar{A}(\Delta) \begin{bmatrix} y\\ \check{z} \end{bmatrix} + \bar{H}(\Delta)[z-\check{z}] + \bar{C}(\Delta)v.$$

The first-order necessary condition for this minimum problem yields

$$v = -\beta \left[-\theta_1 I + \beta \bar{C}(\Delta)' \Omega^*(\Delta^*) \bar{C}(\Delta)\right]^{-1} \bar{C}(\Delta)' \Omega^*(\Delta^*) \left(\bar{A}(\Delta) \begin{bmatrix} y\\ \check{z} \end{bmatrix} + \bar{H}(\Delta) [z - \check{z}] \right)$$

or

$$v = -\bar{F}(\Delta) \begin{bmatrix} z - \check{z} \\ y \\ z \end{bmatrix} = -\bar{F}_1(\Delta)(z - \check{z}) - \bar{F}_2(\Delta) \begin{bmatrix} y \\ z \end{bmatrix}.$$
 (14.2.23)

Equation (14.2.23) gives v, the worst case mean of w^* that comes from applying the \mathbf{T}^1 operator. Conditional on (y, z, \tilde{z}) , the covariance matrix of the worst case distribution of w^* is

$$\Sigma(\Delta) = \left[I - \frac{\beta}{\theta_1} \bar{C}(\Delta)' \Omega^*(\Delta^*) \bar{C}(\Delta)\right]^{-1}.$$

We next compute the matrix $\overline{\Omega}(\Delta)$ in the quadratic form in $[(z - \check{z})' \quad y' \quad \check{z}']$ that emerges from applying the \mathbf{T}^1 operator. First, adjust the objective for the
choice of v by constructing a matrix $\overline{\Pi}(\Delta)$, with row and column dimension both equal to the dimension of $[(z - \tilde{z})' \quad y' \quad \tilde{z}']$:

$$\bar{\Pi}(\Delta) = \begin{bmatrix} 0 & -\tilde{F}_1(\Delta) & \\ I & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{bmatrix}' \begin{bmatrix} Q & P_2 & P_1 & P_2 \\ P_2' & R_{22} - \theta_2 \Delta^{-1} & R_{21} & R_{22} \\ P_1' & R_{21} & R_{11} & R_{12} \\ P_2' & R_{22} & R_{21} & R_{22} \end{bmatrix} \begin{bmatrix} 0 & -\tilde{F}_1(\Delta) & \\ I & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{bmatrix}$$

The matrix in the quadratic form in $[(z - \check{z})' \quad y' \quad \check{z}']$ for the minimized objective function that emerges from applying the \mathbf{T}^1 operator is:

$$\begin{split} \bar{\Omega}(\Delta) &= \bar{\Pi}(\Delta) + \beta \begin{bmatrix} \bar{H}(\Delta)' \\ \bar{A}(\Delta)' \end{bmatrix} \Omega^*(\Delta^*) \Big\{ I + \\ \beta \bar{C}(\Delta) \left[\theta_1 I - \beta \bar{C}(\Omega)' \Omega^*(\Delta^*) \bar{C}(\Delta) \right]^{-1} \bar{C}(\Delta)' \Omega^*(\Delta^*) \Big\} \left[\bar{H}(\Delta) \quad \bar{A}(\Delta) \right]. \end{split}$$

This is a useful formula, as we see next.

14.2.6.1. Worst case distribution for $z - \check{z}$: $\mathcal{N}(u, \Gamma(\Delta))$

Use the partition:⁴

$$\bar{\Omega}(\Delta) = \begin{bmatrix} \bar{\Omega}_{11}(\Delta) & \bar{\Omega}_{12}(\Delta) \\ \bar{\Omega}_{21}(\Delta) & \bar{\Omega}_{22}(\Delta) \end{bmatrix}$$

where $\bar{\Omega}_{11}(\Delta)$ has the same dimension as $z - \check{z}$ and $\bar{\Omega}_{22}(\Delta)$ has the same dimension as $\begin{bmatrix} y \\ \check{z} \end{bmatrix}$. The covariance matrix of $z - \check{z}$ is

$$\Gamma(\Delta) = -\left[\frac{1}{\theta_2}\bar{\Omega}_{11}(\Delta)\right]^{-1},\qquad(14.2.24)$$

which is positive definite when the pair of penalty parameters (θ_1, θ_2) satisfies the no-breakdown condition (14.2.19).

⁴ Knowing $\overline{\Omega}(\Delta)$ allows us to deduce the worst case distribution for $z - \check{z}$ conditional on (y, \check{z}) in another way, thereby establishing a useful cross check on formula (14.2.20) or (14.2.22). See Hansen and Sargent (2005b XXXX).

14.2.7. Worst case signal distribution

Appendix A establishes that the distribution of the signal under the worst case model has conditional mean and conditional covariance matrix:

$$\bar{s}^* = \bar{D}_1(\Delta)y + \bar{D}_2(\Delta)\check{z} + Ha \qquad (14.2.25)$$

$$\bar{\Upsilon} = D_2 \Gamma(\Delta) D_2' + G \Sigma(\Delta) G' \qquad (14.2.26)$$

where

$$\bar{D}_1(\Delta) \doteq D_1 - D_2 \tilde{F}_{21}(\Delta) - G \tilde{F}_{31}(\Delta)$$
$$\bar{D}_2(\Delta) \doteq D_2 - D_2 \tilde{F}_{23}(\Delta) - G \tilde{F}_{33}(\Delta).$$

The laws of motion for (y, \check{z}) under the worst case model are:

$$y^{*} = \left[\Pi_{y}y + \Pi_{s}\bar{D}_{1}(\Delta) - (\Pi_{s}H + \Pi_{a})\tilde{F}_{11}(\Delta)\right]y + \left[\Pi_{s}\bar{D}_{2}(\Delta) - (\Pi_{s}H + \Pi_{a})\tilde{F}_{13}(\Delta)\right]\tilde{z} + \Pi_{s}(s^{*} - \bar{s}^{*}).$$
(14.2.27)

$$\tilde{z}^{*} = \left(A_{21} - B_{2}\tilde{F}_{11}(\Delta) + K_{2}(\Delta)[\bar{D}_{1}(\Delta) - D_{1}]\right)y + \left(A_{22} - B_{2}\tilde{F}_{13}(\Delta) + K_{2}(\Delta)[\bar{D}_{2}(\Delta) - D_{2}]\right)\tilde{z} + K_{2}(\Delta)(s^{*} - \bar{s}^{*}).$$
(14.2.28)
where the innovation $s^{*} - \bar{s}^{*}$ under the distorted model is normal with mean

where the innovation $s^* - \bar{s}^*$ under the distorted model is normal with mean zero and covariance matrix $\bar{\Upsilon}$.

Such a representation for the signal distribution is useful, for example, in asset pricing applications.

14.3. Examples

Example 1 (Jovanovic-Nyarko I):

This example adds adjustment costs to a model authored by Jovanovic and Nyarko (1996XX). A decision maker chooses a scalar decision a_t to maximize

$$E\left[\sum_{t=0}^{\infty} \beta^{t} \left[1 - (s_{t+1} - y_{2t+1})^{2} - da_{t}^{2}\right] |y_{20}\right]$$
(14.3.1)

where

$$y_{2t+1} = y_{2t} + a_t + cw_{t+1} \tag{14.3.2a}$$

$$s_{t+1} = z + gw_{t+1} \tag{14.3.2b}$$

$$z_{t+1} = z_t. (14.3.2c)$$

$$z_0 \sim \mathcal{N}(\check{z}_0, \Delta_0) \tag{14.3.2d}$$

Examples

where w_{t+1} is an i.i.d. 2×1 Gaussian random vector with mean zero and variance I. In (14.3.1), $1 - (s_{t+1} - y_{2t+1})^2$ is the decision maker's time t+1output, and s_{t+1} is an 'ideal' time t+1 action, unknown at t but observed at t+1, that the decision maker aspires to implement by choosing an increment a_t to a prior action y_{2t} , and da_t^2 is the adjustment cost that he pays for altering his prior action. The ideal action is a noisy signal of an unknown mean ideal action z. The decision maker starts with a prior belief \check{z}_0 about the ideal action. Assume that the constant mean z of the ideal action is itself a random variable that is drawn from $\mathcal{N}(\check{z}_0, \Delta_0)$. Assume also that the initial action is $y_0 = \check{z}_0$.

To map this problem into our setting, evaluate $E[(s_{t+1} - y_{2t+1})^2] = E[z - (y_{2t} + a_t)]^2 + (g - c)(g - c)'$, which implies that the time t component of the criterion function is $U(x_t, a_t) = 1 - z^2 - y_{2t}^2 - 2a_ty_{2t} - (d+1)a_t^2 + 2z_ty_{2t} + 2za_t$ plus the variance (g - c)'(g - c). The variance term is beyond control and can be omitted when calculating a decision rule. Let 1 be the first component of the observed part of the state y_t so that

$$\begin{bmatrix} 1\\ y_{2t+1} \end{bmatrix} = \begin{bmatrix} 1 & 0\\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1\\ y_{2,t+1} \end{bmatrix} + \begin{bmatrix} 0\\ 1 \end{bmatrix} a_t + \begin{bmatrix} 0\\ c \end{bmatrix} w_{t+1}$$

which is a version of (14.2.3) with $y_t = \begin{bmatrix} 1 & y_{2t} \end{bmatrix}', A_{11} = I, B_1 = \begin{bmatrix} 0 & 1 \end{bmatrix}', C_1 = \begin{bmatrix} 0 & c \end{bmatrix}', A_{12} = 0$. Comparing (14.3.2b) with (14.2.2c) indicates that we should set $D = \begin{bmatrix} 0 & 1 \end{bmatrix}, H = 0, G = g$, and comparing (14.3.2c) with (14.2.2b) shows that we should set $A_{21} = 0, A_{22} = 1, B_2 = 0, C_2 = 0$. To capture the objective function, we should set $Q = (d+1), P_1 = \begin{bmatrix} 0 & -1 \end{bmatrix}, P_2 = -1, R_{11} = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix}, R_{12} = \begin{bmatrix} 0 \\ -1 \end{bmatrix}, R_{22} = 1, R_{21} = R'_{12}.$

Instructions for Ricardo: This problem is interesting only in the nonstationary case in which Δ_t is moving over time. I ask that you extend your program getStationaryDeltaK2 (i.e., create an alternative program to be used in the nonstationary case) so that starting from a specified Δ_0 to be input by the user, it generates a sequence $\{\Delta_t\}_{t=0}^T$. For many, though not all, problems, Δ_t will converge pretty quickly. Having created $\{\Delta_t\}_{t=0}^T$, then I ask you to modify your game II program so that it kicks out time-varying decision rules for a_t and for the worst case conditional means u_t, \tilde{v}_t, v_t . In addition, could you modify your formulas for the conditional covariances so that they too are time-varying? These are simple modifications for what you already have, implementing as they do the formulas in this chapter or in the HS paper nocommit*.tex. I recommend that you set $c = 1, g = 2, d = 1, \check{z}_0 = 10, \Delta_0 = 5$ to begin. End of instructions for Ricardo.

Example 2: (Jovanovic-Nyarko II):

In example 1, $\lim_{t\to+\infty} \Delta_t = 0$, so that eventually learning about the hidden state is complete. This implies that $\lim_{t\to+\infty} u_t = 0$ and that the associated variance of the distorted distribution for z also converges to zero. To sustain perpetual learning about z and perpetual distortion through a nonzero setting of u, we modify the description of z_t to be:

$$z_{t+1} = (1-\rho)\mu_z + \rho z_t + c_z w_{3,t+1}$$

where $w_{3,t+1} \sim \mathcal{N}(O,1)$ is a third random shock that we add to the model of example 1 and $|\rho| < 1$. The rest of the model remains the same. Now $\lim_{t\to+\infty} \Delta_t = \Delta_{\infty} > 0$, so that there always remains uncertainty about z. In this case, for $\theta_2 < +\infty$, u_t will not converge to zero but instead to a timeinvariant linear function of y_t, \tilde{z}_t that depends on (θ_1, θ_2) .

Example 3 (pure estimation of hidden state):

Suppose that B = 0 so that the state cannot be influenced by a. The decision maker wants to estimate -Px. Set Q = I and R = P'P. The decision a_t as a function of (y_t, \tilde{z}_t) is a robust estimate of $-Px_t$.

Example 4 (Hansen, Sargent, and Wang):

In chapter 12,⁵ we used a model of Hansen, Sargent, and Tallarini's (HST, 1999) to describe some of the implications of a preference for robustness for asset pricing. Hansen, Sargent, and Wang (HSW, 2002) modified HST's permanent income model by withholding knowledge of the two separate components of the endowment process posited by HST, thereby impelling the representative to base consumption-saving decisions on filtered estimates of those two components. HSW used their model to study the effects of filtering on market prices of risk. HSW's representative agent faces a problem that falls within the setting of this chapter.

⁵ Programs that do the calculations are in C:/tom/projects/robust/games/hst/forricardom and one directory up. The general HST programs are in c:/tom/projects/hansen/hst.

Examples

We briefly summarize HST's model, which was described in detail in chapters 9 and 12. A planner values a scalar process s of consumption services according to

$$V_0 = -\sum_{t=0}^{\infty} \beta^t (s_t - b_t)^2$$
(14.3.3)

where the service s is produced by the scalar consumption process c via the household technology

$$s_{t} = (1 + \lambda)c_{t} - \lambda h_{t-1}$$

$$h_{t} = \delta_{h}h_{t-1} + (1 - \delta_{h})c_{t}$$
(14.3.4)

where $\lambda \geq 0$ and $\delta_h \in (0,1)$, b is a constant bliss level of services, and h is a scalar stock of household habits.

There is a linear technology for converting an exogenous scalar endowment d_t into consumption or capital:

$$c_t + k_t = Rk_{t-1} + d_t \tag{14.3.5}$$

where k_t, d_t are the capital stock and the exogenous stochastic endowment at time t, respectively. HST showed that R is the gross return on the risk free asset, made constant by the technology.

HST assumed that the agent and the planner observe histories of each component of the following two-component model for the endowment:⁶

$$\begin{split} d_{t+1} &= \mu_d + d_{t+1}^1 + d_{t+1}^2 \\ d_{t+1}^1 &= g_1 d_t^1 + g_2 d_{t-1}^1 + c_1 \epsilon_{t+1}^1 \\ d_{t+1}^2 &= a_1 d_t^2 + a_2 d_{t-1}^2 + c_2 \epsilon_{t+1}^2, \end{split}$$

where $\epsilon_{t+1} = \begin{bmatrix} \epsilon_{t+1}^1 \\ \epsilon_{t+1}^2 \end{bmatrix}$ is an i.i.d. 2×1 Gaussian disturbance vector with mean zero and identity covariance matrix.

HST assumed that the planner observes current and lagged values of both components d_t^i at time t. HSW assumed instead that the planner sees only current and lagged values of the sum d_t at time t.

⁶ Tom: Note that in the earlier chapter and HST we instead used $d^1 = d^*, d^2 = \hat{d}, c^1 = c^*, c^2 = \hat{c}$. Change this in the text.

The law of motion for the state of the model can be written

$$\begin{bmatrix} h_t \\ k_t \\ d_t \\ 1 \\ d_{t+1} \\ d_t^{1+1} \\ d_t^{1+1} \\ d_t^{1+1} \end{bmatrix} = \begin{bmatrix} \delta_h & (1-\delta_h)\gamma & 0 & 0 & (1-\delta_h) & 0 & 0 \\ 0 & \delta_k & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & g_1 & g_2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} h_{t-1} \\ k_{t-1} \\ d_{t-1} \\ 1 \\ d_t \\ d_t^{1} \\ d_t^{1} \end{bmatrix}$$

$$+ \begin{bmatrix} -(1-\delta_h) \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} i_t + \begin{bmatrix} 0 & 0 \\ 0 &$$

This fits our general setting with

$$y_t = \begin{bmatrix} h_{t-1} \\ k_{t-1} \\ d_{t-1} \\ 1 \\ d_t \end{bmatrix}, \quad z_t = \begin{bmatrix} d_t^1 \\ d_{t-1}^1 \\ \end{bmatrix}, \quad s_t = y_t.$$

For HST, the planner knows current and lagged values of y_t and z_t when i_t is chosen. HSW instead assume that current and lagged values of y_t are in the planner's information set when i_t is to be chosen, but that z_t is never observed. The planner bases his decisions on an estimate of z_t gleaned from the history of y_t . This makes the Bellman equation for the robust planner a version of (14.2.18).

Lars and Tom: We should generalize the observational equivalence result of HST, and the extension in the earlier chapter, to the two- θ case.

14.3.1. Asset pricing with robustness and filtering

As noted in section 14.2.5, although game (14.2.12) takes account of the fact that part of the state is estimated, in the end it still takes the form of the robust linear regulator without filtering of the kind analyzed in chapters 2 and 6. The asset pricing theory from chapter 12 applies directly. HSW used this theory to construct multi-period versions of market prices of risk for their model and compared them with HST's model.

A. Worst case signal distribution

In this appendix, we construct a recursive representation of the distribution of signals under the distorted probability distribution. Recall the signal evolution:

$$s^* = Dx + Ha + Gw^*.$$

Under the approximating model, the signal next period is normal with mean

$$\check{s}^* = D_1 y + D_2 \check{z} + Ha$$

and covariance matrix

$$\check{\Upsilon} = D_2 \Delta D_2' + GG'.$$

The distorted mean of the signal conditioned on the signal history is:

$$\bar{s}^* = D_1 y + D_2 \check{z} + (D_2 u + G \tilde{v}) + Ha$$

which by virtue of the second and third blocks of rows of (14.2.22) can be written

$$\bar{s}^* = \bar{D}_1(\Delta)y + \bar{D}_2(\Delta)\check{z} + Ha \tag{14.A.1}$$

where

$$\bar{D}_1(\Delta) \doteq D_1 - D_2 \tilde{F}_{21}(\Delta) - G \tilde{F}_{31}(\Delta)$$
$$\bar{D}_2(\Delta) \doteq D_2 - D_2 \tilde{F}_{23}(\Delta) - G \tilde{F}_{33}(\Delta).$$

The distorted covariance matrix is:

$$\bar{\Upsilon} = D_2 \Gamma(\Delta) D_2' + G \Sigma(\Delta) G'.$$

To construct the distorted dynamics for y^* , start from the formula $y^* = \prod_s s^* + \prod_y y + \prod_a a$. Substituting for the robust decision rule for a from the first block of row of (14.2.22) and replacing s^* with with $\bar{s}^* + (s^* - \bar{s}^*)$ from (14.A.1) gives

$$y^{*} = [\Pi_{y}y + \Pi_{s}\bar{D}_{1}(\Delta) - (\Pi_{s}H + \Pi_{a})\tilde{F}_{11}(\Delta)]y + [\Pi_{s}\bar{D}_{2}(\Delta) - (\Pi_{s}H + \Pi_{a})\tilde{F}_{13}(\Delta)]\check{z} + \Pi_{s}(s^{*} - \bar{s}^{*}).$$
(14.A.2)

To complete a recursive representation for y^* under the worst case distribution, we need a formula for updating \check{z}^* under the worst case distribution. Recall the formula for \check{z}^* under the approximating model from the Kalman filter (14.2.6b) or (14.2.7):

$$\check{z}^* = [A_{21} - B_2 \tilde{F}_{11}(\Delta)]y + [A_{22} - B_2 \tilde{F}_{13}(\Delta)]\check{z} + K_2(\Delta)(s^* - D_1 y - D_2 \check{z} - Ha)$$

or

$$\check{z}^* = [A_{21} - B_2 \tilde{F}_{11}(\Delta)]y + [A_{22} - B_2 \tilde{F}_{13}(\Delta)]\check{z} + K_2(\Delta)(s^* - \check{s}^*).$$

Using the identity

$$s^* - \check{s}^* = (s^* - \bar{s}^*) + (\bar{s}^* - \check{s}^*)$$

= $(s^* - \bar{s}^*) + ([\bar{D}_1(\Delta) - D_1]y + [\bar{D}_2(\Delta) - D_2]\check{z})$

in the above equation gives:

$$\tilde{z}^{*} = \left(A_{21} - B_{2}\tilde{F}_{11}(\Delta) + K_{2}(\Delta)[\bar{D}_{1}(\Delta) - D_{1}]\right)y
+ \left(A_{22} - B_{2}\tilde{F}_{13}(\Delta) + K_{2}(\Delta)[\bar{D}_{2}(\Delta) - D_{2}]\right)\check{z} + K_{2}(\Delta)(s^{*} - \bar{s}^{*}).$$
(14.A.3)

Taken together, (14.A.2) and (14.A.3) show how to construct \check{z}^* from the signal history under the distorted law of motion. The innovation $s^* - \bar{s}^*$ under the distorted model is normal with mean zero and covariance matrix $\check{\Upsilon}$.

Part V More applications

Chapter 15 Multiple agents

15.1. Introduction

This chapter and the next describe equilibria in which multiple decision makers share an approximating model and are all concerned about model misspecification. Because we desire to preserve as much as possible of the structure and empirical power of rational expectations, we impute a common approximating model to all of the agents When they have different objective functions, the context-specific worst case models of different decision makers will differ. We thus have a highly structured way of modelling what seem to be *ex post* heterogeneous beliefs, a point we discuss further in subsection 15.2.2.

In the present chapter, we study two player dynamic games. We adapt the concept of Markov perfect equilibrium to incorporate concerns about robustness to model misspecification. Here the timing protocol is that both players choose sequentially and simultaneously. In chapter 16, we study a different timing protocol in which a Stackelberg leader chooses once and for all at time 0, while Stackelberg followers choose sequentially.

15.2. Markov perfect equilibria with robustness

The decisions of two agents affect the motion of a common state vector that impinges on the return functions of both agents. Without concerns about robustness, a Markov perfect equilibrium can be computed by working backwards on pairs of Bellman functions and the associated equations that express decision rules as functions of continuation value functions.¹ Similar procedures apply when we impute concerns about robustness to both decision makers. Our equilibrium concept insists that the decision makers share an approximating model.

 $^{^1\,}$ See Ljungqvist and Sargent (2004,XXXX chapter 7).

⁻³⁵⁹ -

For each agent, the approximating model incorporates the robust decision rule used by the other agent.

The model is

$$x_{t+1} = Ax_t + B_1 u_{1t} + B_2 u_{2t} + C\epsilon_{t+1}$$
(15.2.1)

where u_{it} is a control vector chosen by agent *i* as a function of the state x_t , and ϵ_{t+1} is an i.i.d. Gaussian random vector with mean zero and identity covariance matrix. Agent *i* conceives of model misspecification by thinking that the actual data generating mechanism comes from a set of perturbations to (15.2.1) of the form

$$x_{t+1} = Ax_t + B_1 u_{1t} + B_2 u_{2t} + C(\epsilon_{t+1} + w_{it+1})$$
(15.2.2)

where w_{it+1} represents misspecified dynamic components that depend on the history of x_s up to time t. Agent i wants to maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t r_i(x_t, u_{it})$$
 (15.2.3)

where $\beta \in (0,1)$ and $r_i(x_t, u_{it}) = -[x'_t R_i x_t + u'_{it} Q_i u_{it} + 2u'_{it} H_i x_t].$

We appeal to the version of certainty equivalence cited on page 27 to allow us to drop the ϵ_{t+1} term from (15.2.2) and the conditional expectation E from (15.2.3) and proceed to solve nonstochastic versions of both players' extremum problems.

We define a Nash equilibrium with robust decision makers and a common approximating model. In equilibrium, player i selects a robust decision rule of the form

$$u_{it} = -F_{it}x_t.$$
 (15.2.4)

Though in the limit we will seek a time invariant rule F_i , to accommodate backward induction we begin by allowing time-varying rules. The set of laws of motion confronting agent i has the form

$$x_{t+1} = (A - B_{-i}F_{-it})x_t + B_iu_{it} + Cw_{it+1}$$
(15.2.5)

where a subscript $_{-i}$ refers to the other player. Notice that (15.2.5) incorporates the robust rule F_{-it} of the other player and that each player has his own distortion process w_{it} . Player *i* solves a multiplier robust control problem with multiplier θ_i .

Definition 15.2.1. A Markov perfect equilibrium with concerns about robustness consists of pairs of value functions V_i , decision rules $u_i = -F_i x_i$, and rules for worst case shocks $w_i = K_i x_i$ such that the decision rules for u_i, w_i attain $V_i(x)$ and the value functions V_i satisfy the Bellman equations

$$V_i(x) = \max_{u_i} \min_{w_i} \{ r_i(x, u_i) + \beta \theta_i w'_i w_i + \beta V_i(x^*) \}$$
(15.2.6)

where * denotes next period's value and the extremization is subject to

$$x^* = (A - B_{-i}F_{-i})x + B_iu_i + Cw_i.$$
(15.2.7)

The value functions assume the form

$$V_i(x) = -x'P_ix_i$$

where $P_i = T_i \circ \mathcal{D}_i P_i$ is a fixed point defined in terms of the composition of modified versions of two familiar operators:

$$T_{i}(P_{i}) = Q_{i} + \beta (A - B_{-i}F_{-i})'P_{i}(A - B_{-i}F_{-i}) - (\beta (A - B_{-i}F_{-i})'P_{i}B_{i} + H_{i}')(R_{i} + \beta B_{i}'P_{i}B_{i})^{-1} \times (\beta B_{i}'P_{i}(A - B_{-i}F_{-i}) + H_{i})$$
(15.2.8)

$$\mathcal{D}_i(P_i) = P_i + \theta_i^{-1} P_i C (I - \theta^{-1} C P_i C)^{-1} C' P_i.$$
(15.2.9)

The T_i operator is associated with the maximization part of the problem on the right side of (15.2.7), while the \mathcal{D}_i operator is associated with the minimization part.

In the next subsection, we describe a recursive algorithm for computing a Markov perfect equilibrium with concerns about robustness.

15.2.1. Computational algorithm: iterating on stacked Bellman equations

Define the iterations

$$F_{it} = (R_i + \beta B'_i \mathcal{D}_i (P_{it+1}) B_i)^{-1} (\beta B'_i P_{it+1} (A - B_{-i} F_{-it}) + H) (15.2.10)$$

$$P_{it} = T_i \circ \mathcal{D}_i (P_{it+1}).$$
(15.2.11)

We propose to use iterations on these operators to find fixed points $F_i, P_i, i = 1, 2$, that satisfy

$$F_i = (R_i + \beta B'_i \mathcal{D}_i(P_i) B_i)^{-1} (\beta B'_i P_i (A - B_{-i} F_{-i}) + H)$$
(15.2.12)

$$P_i = T_i \circ \mathcal{D}_i(P_i). \tag{15.2.13}$$

Suppose that u_i is $k_i \times n$. Given P_{1t+1}, P_{2t+1} , equations (15.2.10) for i = 1, 2 form $(k_1 + k_2) \times n$ linear equations in the same number of variables, namely, F_{1t}, F_{2t} . To compute an equilibrium, start with zero terminal value matrices P_{1T}, P_{2T} , solve (15.2.10) for F_{1T}, F_{2T} , then iterate backwards on (15.2.10),(15.2.11) until, hopefully, the F_{it}, P_{it} sequences converge. If they converge, we say that there is an asymptotically time invariant equilibrium law of motion.

When both players use time invariant robust rules, the approximating model becomes

$$x_{t+1} = A^* x_t + C\epsilon_{t+1} \tag{15.2.14}$$

where $A^* = A - B_1F_1 - B_2F_2$ and where we have reactivated the Gaussian disturbance. The two agents share this approximating model but in general have different worst case models. The worst case model for agent *i* is

$$x_{t+1} = A^* x_t + C(\epsilon_{t+1} + w_{it+1})$$

$$w_{it+1} = K_i x_t$$

where

$$K_i = \theta^{-1} (I - \theta_i^{-1} C' P_i C)^{-1} C' P_i A^*.$$
(15.2.15)

Another expression for the worst case model of player i is

$$x_{t+1} = (A^* + CK_i)x_t + C\epsilon_{t+1}.$$
(15.2.16)

15.2.2. Bayesian interpretation and belief heterogeneity

A version of our usual 'ex post Bayesian interpretation' of each player's robust rule applies. After we have computed an equilibrium and know the different worst-case shocks $w_{it} = K_i x_t$ of the two players, each player *i* can be regarded as solving an ordinary control problem, using its own twisted law of motion (15.2.16) and taking as given the decision rule $u_{-i,t} = -F_{-i}x_t$ of the other player. Notice how this builds in full understanding about the other player's decision rule, a counterpart to a rational expectations assumption.

Thus, we have a highly disciplined way of taking a common approximating model and from it generating what appear to be heterogeneous beliefs. The fact that Bellman equations for the zero-sum two-player game solved by each player are the sources of those heterogeneous beliefs means that cross-equation restrictions very similar to those characterizing rational expectations models (e.g., see Hansen and Sargent (1980, 1981 XXX)) will still prevail.

15.3. Concluding remarks

We have proposed an equilibrium concept in which both players in a dynamic game share a common approximating model but are concerned that it is misspecified. Their disparate motives means that their worst-case beliefs twist the approximating model in different directions. We have shown how this kind of equilibrium can be computed by adapting existing methods for computing Markov perfect equilibria without concerns about robustness. In the next chapter, we apply a similar equilibrium concept to a setting with a timing protocol that requires one agent to commit and the other agents to choose sequentially.

Chapter 16 Robustness in forward looking models

16.1. Introduction

This chapter continues the enterprise begun in chapter 15 of studying situations in which differently situated and differently motivated agents all share a common approximating model and want robust decision rules.¹ In the interests of formulating robust versions of optimal policy problems, we alter the timing protocols from those in chapter 15 in order to incorporate a structure with a 'leader' who must commit to future contingency plans and some 'followers' who must choose sequentially. Here and in chapter 18 we study Stackelberg plans constructed with an eye to concerns about misspecification on the parts of both the Stackelberg leader and the followers. For example, in chapter 18, a benevolent government acts as a Stackelberg leader with respect to a competitive private sector, modelled in terms of a representative consumer who acts as a follower. At time 0, the leader chooses a sequence of actions, taking into account how a follower's decisions at each date will respond to its forecasts of future actions by the leader. The leader's policy instruments appear as 'forcing variables' in a follower's Euler equations. Those Euler equations describe how a follower's decisions depend on the sequence of the leader's actions. When a follower is concerned about robustness, some of those Euler equations describe the follower's worst case shocks, as we saw in chapter 6.

Without concerns about robustness, a 'first-order' approach to solving Stackelberg problems is to use the followers' Euler equations to summarize their best responses to the leader's decisions, then to form a Lagrangian for the leader with a sequence of multipliers adhering to the followers' Euler equations. The followers' Euler equations are 'implementability constraints' that require the leader's decision at time t to confirm forecasts that a follower anticipated when he made his earlier decisions. The Lagrange multipliers on the implementability constraints make the leader's actions depend on the history of outcomes and allow a recursive representation of the leader's decision rule.

 $^{^{1}}$ This chapter builds upon and corrects aspects of Hansen and Sargent (2003).

⁻³⁶⁵ -

This chapter formulates an equilibrium in which the leader and the followers are both concerned about model misspecification. As a natural counterpart of rational expectations, we assume that the leader and followers share a common approximating model. When both types of agent are concerned about robustness, the approximating model for one agent must include a description of the robust decision rules of the other types of agent, and of how they respond to other decision makers' actions.

How can the leader understand how he affects the robust decision rules of the followers? The leader can influence the followers' worst case shocks and thereby their decision rules. The essential insight in this chapter is to note that we can appeal to results in chapter 6 to assert that the followers' Euler equations, including those for the worst case shocks, fully summarize the robust followers' best response to the leader; and that by augmenting the followers' Euler equations to include those for their worst-case shocks, a standard method for computing a Ramsey plan can be applied.

The remainder of this chapter is organized as follows. Section 16.2 states a Stackelberg problem in which decision makers fear model misspecification and therefore want robustness. Section 16.3 describes how to solve the robust Stackelberg problem by properly rearranging and reinterpreting some state variables and some Lagrange multipliers after having solved a robust linear regulator problem. As an example, section 16.5 describes a dynamic model of a monopolist facing a competitive fringe. Section 16.6 concludes. Appendix A describes how the invariant subspace methods of chapter 3 can also be used to compute robust Ramsey plans. Appendix B studies the Riccati equation that solves the robust Ramsey problem. Appendix C describes the connection of our work to a Bellman equation that Marcet and Marimon (1999) have used to solve problems with implementability constraints like ours.

16.1.1. Related literature

Brunner and Meltzer (1969) and Von Zur Muehlen (1982) were early advocates of zero-sum two person games for representing model uncertainty and designing macroeconomic rules.² Stock (1999), Sargent (1999), and Onatski and Stock (1999) have used versions of robust control theory to study robustness of purely backward looking macroeconomic models. They focused on whether a concern for robustness would make policy rules more or less aggressive in response to shocks. Vaughan (1970), Blanchard and Khan (1980), Whiteman (1983), and Anderson and Moore (1985) are early sources on solving control problems with forward-looking private sectors.³ Without a concern for robustness, Kydland and Prescott (1980), Hansen, Epple, and Roberds (1985), Miller and Salmon (1985a, 1985b), Backus and Driffill (1986), Sargent (1987), Currie and Levine (1987), Pearlman, Currie, and Levine (1986), Pearlman (1992), Woodford (1998), King and Wolman (1999), and Marcet and Marimon (1999) have solved Stackelberg or Ramsey problems using Lagrangian formulations. Pearlman, Currie, and Levine (1986), Pearlman (1992) and Svensson and Woodford (2000) study the control of forward looking models where part of the state is unknown and must be filtered. DeJong, Ingram, and Whiteman (1996), Otrok (XXXX), and others study the Bayesian estimation of forward looking models. They summarize the econometrician's doubts about parameter values with a prior distribution, meanwhile attributing no doubts about parameter values to the private agents in their models. Mark Giannoni (2000) studies robustness in a forward looking macro model. He models the policy maker as knowing all parameters except two, for which he knows only bounds. The policy maker then computes the $\min - \max$ policy rule. Kasa (2000) also studies robust policy in a forward looking model. Onatski (2000) designs simple (not history dependent) robust policy rules for a forward looking monetary model. Christiano and Gust (1999) study robustness from the viewpoint of the determinacy and stability of rules under nearby parameters. They adopt a perspective of robust control theorists like Basar and Bernhard (1995) and Zhou, Doyle, and Glover (1996), who are interested in finding rules that stabilize a system under the largest possible set of departures from a reference model.

 $^{^2\,}$ More generally, Hurwicz (1951) had advocated zero-sum games as a way of making decisions when a decision maker could not specify a unique model.

³ Chapter 3 describes efficient computational algorithms for such models.

Relative to the literature, a novelty in the approach of this chapter is that we impute concerns about model misspecification to both the leader and the follower. While some of the papers that we have just cited attribute model uncertainty to the leader (a.k.a. the government), all of them assume that the followers have no doubts about the model.

16.2. The robust Stackelberg problem

This section defines a robust Stackelberg problem where the Stackelberg leader is concerned about model misspecification. In macroeconomic problems, the Stackelberg leader is often a government and the Stackelberg follower is a representative agent within a private sector. In section 16.5, we present an application with a monopolist and a competitive fringe.

Let z_t be an $n_z \times 1$ vector of natural state variables, x_t an $n_x \times 1$ vector of endogenous variables free to jump at t, and U_t a vector of the leader's instruments. The z_t vector is inherited from the past. The model determines the 'jump variables' x_t at time t. Included in x_t are prices and quantities that adjust to clear markets at time t. Let $y_t = \begin{bmatrix} z_t \\ x_t \end{bmatrix}$. Define the Stackelberg leader's one-period loss function⁴

$$r(y,u) = y'Qy + u'Ru.$$
 (16.2.1)

The leader wants to maximize

$$-\sum_{t=0}^{\infty} \beta^t r(y_t, U_t).$$
 (16.2.2)

The leader makes policy in light of a set of models indexed by a vector of specification errors W_{t+1} around its approximating model:

$$\begin{bmatrix} I & 0 \\ G_{21} & G_{22} \end{bmatrix} \begin{bmatrix} z_{t+1} \\ x_{t+1} \end{bmatrix} = \begin{bmatrix} \hat{A}_{11} & \hat{A}_{12} \\ \hat{A}_{21} & \hat{A}_{22} \end{bmatrix} \begin{bmatrix} z_t \\ x_t \end{bmatrix} + \hat{B}U_t + \hat{C}W_{t+1}.$$
(16.2.3)

⁴ The problem assumes that there are no cross products between states and controls in the return function. A simple transformation converts a problem whose return function has cross products into an equivalent problem that has no cross products. See Chapter 3. XXXX Add page tag.

We assume that the matrix on the left is invertible, so that⁵

$$\begin{bmatrix} z_{t+1} \\ x_{t+1} \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} z_t \\ x_t \end{bmatrix} + BU_t + CW_{t+1}$$
(16.2.4)

or

$$y_{t+1} = Ay_t + BU_t + CW_{t+1}.$$
(16.2.5)

The followers' behavior is summarized by the second block of equations of (16.2.3) or (16.2.4). These typically include the first-order conditions of private agent's optimization problem (i.e., their Euler equations). These equations summarize the forward looking aspect of the followers' behavior. The particular structure of these equations and the variables composing x_t depend on the followers' optimization problems, and in particular, whether we impute a concern about robustness to them. As we shall see later, if we want to impute a motive for robustness to the followers, then we must include w_{t+1} , the specification errors of the followers, among the variables in x_t , and we must include Euler equations pertaining to the choice of w_{t+1} among the second block of equations of (16.2.3) and (16.2.4). In section 16.5, we'll analyze an example.

Returning to (16.2.3) or (16.2.4), the vector W_{t+1} of unknown specification errors can feed back, possibly nonlinearly, on the history y^t , which lets the W_{t+1} sequence represent misspecified dynamics in the leader's approximating model. The leader regards its approximating model (which asserts that $W_{t+1} = 0$) as a good approximation to the unknown true model in the sense that the unknown W_{t+1} sequence satisfies⁶

$$\sum_{t=0}^{\infty} \beta^{t+1} W_{t+1}' W_{t+1} \le \eta_0 \tag{16.2.6}$$

where $\eta_0 > 0$.

The certainty equivalence principle stated on page 27 allows us to work with non stochastic approximating and distorted models. We would attain the same decision rule if we were to replace x_{t+1} with the forecast $E_t x_{t+1}$ and to

 $^{^{5}}$ We have assumed that the matrix on the left of (16.2.3) is invertible for ease of presentation. However, by appropriately using the invariant subspace methods described in Chapter 3 and appendix A, it is straightforward to adapt the computational method when this assumption is violated.

⁶ If $C\epsilon_{t+1}$ were added to the right side of (16.2.5), we would take the expectation of (16.2.6).

add a shock process $C\epsilon_{t+1}$ to the right side of (16.2.3) or $C\epsilon_{t+1}$ to the right side of (16.2.4), where ϵ_{t+1} is an i.i.d. random vector with mean of zero and identity covariance matrix.

Let X^t denote the history of any variable X from 0 to t. Kydland and Prescott (1980), Miller and Salmon (1982, 1985), Hansen, Epple, and Roberds (1985), Pearlman, Currie and Levine (1986), Sargent (1987), Pearlman (1992) and others have studied non-robust (i.e., $\eta_0 = 0$) versions of the following problem:

Definition 16.2.1. For $\eta > 0$, the constraint version of the Stackelberg or Ramsey problem is to extremize (16.2.2) by finding a sequence of decision rules expressing U_t and W_{t+1} as sequences of functions mapping the time t history of the state z^t into the time t decision. The leader chooses these decision rules at time 0 and commits to them everyone.

Definition 16.2.2. When $\eta_0 > 0$, the decision rule for U_t that solves the Stackelberg problem is called a robust *Stackelberg plan* or robust *Ramsey plan*.

Note that the decision rules are designed to depend on the history of the true state z_t and not on the history of the jump variable x_t . For a non-robust version of the problem, the aforementioned authors show that the optimal rule is history-dependent, which in our context means that U_t, W_{t+1} depend not only on z_t but also on its lags. The history dependence comes from two sources: (a) the leader's ability to commit to a sequence of rules at time $0,^7$ and (b) the forward-looking behavior of the followers that is embedded in the second block of equations in (16.2.3) or (16.2.4).

Fortunately, there is a recursive way of expressing this history dependence by having decisions U_t, W_{t+1} depend linearly on the current value z_t and on μ_{xt} , a vector of Lagrange multipliers on the last n_x equations of (16.2.3) or (16.2.4). A solution of the problem in Definition 16.2.2 implies a law of motion that expresses μ_{xt+1} as a linear function of (z_t, μ_{xt}) . The dynamics of μ_{xt} express the history dependence of the leader's plan. These multipliers track past leader promises about current and future settings of U. If at time 0 there are no past promises to honor, it is appropriate for the leader to initialize the multipliers to zero (this maximizes its criterion function). The multipliers

⁷ The leader would make different choices were it to choose sequentially, that is, were it to set U_t at time t rather than at time 0.

take non-zero values thereafter, reflecting the subsequent costs to the leader of adhering to its time 0 commitments.

16.2.1. Multiplier version of the robust Stackelberg problem

In chapter 6 and 7, we showed that it is usually more convenient to solve a multiplier game rather than a constraint game. Accordingly, we use:

Definition 16.2.3. The multiplier version of the robust Stackelberg problem is the zero-sum two-player game:

$$\max_{\{U_t\}_{t=0}^{\infty}} \min_{\{W_{t+1}\}_{t=0}^{\infty}} - \sum_{t=0}^{\infty} \beta^t \left\{ r(y_t, U_t) - \beta \Theta W'_{t+1} W_{t+1} \right\}$$
(16.2.7)

where the extremization is subject to (16.2.5) and $\underline{\Theta} < \Theta < \infty$.

16.3. Solving the robust Stackelberg problem

This section describes a three step algorithm for solving a multiplier version of a robust Stackelberg problem.

16.3.1. Step 1: solve a robust linear regulator

Step 1 temporarily disregards the forward looking aspect of the problem (step 3 will take account of that) and notes that superficially the multiplier version of the robust Stackelberg problem (16.2.7), (16.2.5) has the form of a robust linear regulator problem. Mechanically, we can solve this artificial robust linear regulator by noting that associated with problem (16.2.7) is the Bellman equation⁸

$$v(y) = \max_{u} \min_{W} \left\{ -r(y, u) + \beta \Theta W' W + \beta v(y^*) \right\},$$
 (16.3.1)

 $^{^{8}}$ By following the approaches of Kydland and Prescott (1980) and Marcet and Marimon (2000), appendix C describes a closely related Bellman equation that can be used to compute a robust Ramsey plan.

where y^* denotes next period's value of the state and the extremization is subject to the transition law $y^* = Ay + Bu + CW$. The value function that satisfies (16.3.1) has the form v(y) = -y'Py, where P is a fixed point of the operator $T \circ \mathcal{D}$ defined in chapters 2 and 6, namely,

$$T(P) = Q + \beta A' P A - \beta^2 A' P B (R + \beta B' P B)^{-1} B' P A$$
(16.3.2)

$$\mathcal{D}(P) = P + \Theta^{-1} P C (I - \Theta^{-1} C' P C)^{-1} C' P.$$
(16.3.3)

Thus, the Bellman equation (16.3.1) leads to the Riccati equation

$$P = T \circ \mathcal{D}(P). \tag{16.3.4}$$

The *T* operator emerges from the maximization over *U* on the right side of (16.3.1), while the \mathcal{D} operator emerges from the minimization over *W*. The extremizing decision rules are given by $U_t = -F_1 y_t$ where

$$F = \beta (R + \beta B' \mathcal{D}(P)B)^{-1} B' \mathcal{D}(P)A$$
(16.3.5)

and $W_{t+1} = Ky_t$ where

$$K = \Theta^{-1} (I - \Theta^{-1} C' P C)^{-1} C' P (A - BF).$$
(16.3.6)

(See page 29.) Everything that we need to solve the robust Stackelberg problem is encoded in the triple (P, F, K). The next steps decode the solution of the Riccati equation $P = T \circ \mathcal{D}$ to recover objects that solve the robust Stackelberg problem.

16.3.2. Step 2: use the stabilizing properties of shadow price Py_t

We use P to describe how shadow prices on the transition law relate to the artificial state vector $y_t = [z'_t \ x'_t]'$ (we say 'artificial' because x_t is a vector of jump variables.) Recall the Lagrangian methods used in chapters 3 and 6. Thus, another way to solve the multiplier version of the robust Stackelberg problem (16.2.7), (16.2.5) is to form the Lagrangian:

$$\mathcal{L} = -\sum_{t=0}^{\infty} \beta^t \Big[y_t' Q y_t + U_t' R U_t + 2\beta \mu_{t+1}' (A y_t + B U_t + C W_{t+1} - y_{t+1}) - \beta \Theta W_{t+1}' W_{t+1} \Big].$$
(16.3.7)

We want to maximize (16.3.7) with respect to sequences for U_t and y_{t+1} and minimize it with respect to a sequence for W_{t+1} . The first-order conditions with respect to U_t, y_t, W_{t+1} , respectively, are:

$$0 = RU_t + \beta B' \mu_{t+1} \tag{16.3.8a}$$

$$\mu_t = Qy_t + \beta A' \mu_{t+1} \tag{16.3.8b}$$

$$0 = \beta \Theta W_{t+1} - \beta C' \mu_{t+1}. \tag{16.3.8c}$$

Solving (16.3.8*a*) and (16.3.8*c*) for U_t and W_{t+1} and substituting into (16.2.5) gives

$$y_{t+1} = Ay_t - \beta (BR^{-1}B' - \beta^{-1}\Theta^{-1}CC')\mu_{t+1}.$$
 (16.3.9)

Write (16.3.9) as

$$y_{t+1} = Ay_t - \beta \tilde{B} \tilde{R}^{-1} \tilde{B}' \mu_{t+1}.$$
 (16.3.10)

We can represent the system formed by (16.3.10) and (16.3.8b) as

$$\begin{bmatrix} I & \beta \tilde{B} \tilde{R}^{-1} \tilde{B}' \\ 0 & \beta A' \end{bmatrix} \begin{bmatrix} y_{t+1} \\ \mu_{t+1} \end{bmatrix} = \begin{bmatrix} A & 0 \\ -Q & I \end{bmatrix} \begin{bmatrix} y_t \\ \mu_t \end{bmatrix}$$
(16.3.11)

or

$$L^* \begin{bmatrix} y_{t+1} \\ \mu_{t+1} \end{bmatrix} = N \begin{bmatrix} y_t \\ \mu_t \end{bmatrix}.$$
(16.3.12)

We want to find a stabilizing solution of (16.3.12), i.e., one that satisfies

$$\sum_{t=0}^{\infty} \beta^t y_t' y_t < +\infty.$$

The stabilizing solution is attained by setting $\mu_0 = Py_0$, where P solves the matrix Riccati equation $P = T \circ \mathcal{D}(P)$. The solution $\mu_0 = Py_0$ replicates itself over time in the sense that

$$\mu_t = P y_t. \tag{16.3.13}$$

16.3.3. Key insight

In a typical robust linear regulator problem, y_0 is a state vector inherited from the past; the multiplier μ_0 jumps at t = 0 to satisfy $\mu_0 = Py_0$. See chapter 3. But in the Stackelberg problem, pertinent components of both y_0 and μ_0 must adjust to satisfy $\mu_0 = Py_0$, as shown in step 3.

16.3.4. Step 3: convert implementation multipliers into state variables

Partition μ_t conformably with the partition of y_t into $\begin{bmatrix} z'_t & x'_t \end{bmatrix}'$:⁹

$$\mu_t = \begin{bmatrix} \mu_{zt} \\ \mu_{xt} \end{bmatrix}.$$

For the robust Stackelberg problem, only the first n_z elements z_t of $y_t = [z_t \ x_t]'$ are predetermined and the remaining x_t components are free. And while the first n_z elements μ_{zt} of μ_t are free to jump at t, the remaining components μ_{xt} are not. The third step completes the solution of the robust Stackelberg problem by taking note of these facts. We convert the last n_x Lagrange multipliers μ_{xt} into state variables by using the following procedure after we have performed the key step of computing the P that solves the Riccati equation $P = T \circ \mathcal{D}(P)$.

Write the last n_x equations of (16.3.13) as

$$\mu_{xt} = P_{21}z_t + P_{22}x_t. \tag{16.3.14}$$

⁹ This argument simply adapts one in Pearlman (1992). The Lagrangian associated with the robust Stackelberg problem remains (16.3.7). Then the logic of section 16.3.2 implies that the stabilizing solution must satisfy (16.3.13). It is only in how we impose (16.3.13) that the solution diverges from that for the linear regulator.

The vector μ_{xt} becomes part of the state at t, while x_t is free to jump at t. Therefore, solve (16.3.13) for x_t in terms of (z_t, μ_{xt}) :

$$x_t = -P_{22}^{-1}P_{21}z_t + P_{22}^{-1}\mu_{xt}.$$
 (16.3.15)

Then we can write

$$y_t \equiv \begin{bmatrix} z_t \\ x_t \end{bmatrix} = \begin{bmatrix} I & 0 \\ -P_{22}^{-1}P_{21} & P_{22}^{-1} \end{bmatrix} \begin{bmatrix} z_t \\ \mu_{xt} \end{bmatrix}.$$
 (16.3.16)

With these modifications, the key formulas (6.10.2) and (16.3.4) from the optimal linear regulator for F and P, respectively, continue to apply. Using (16.3.16), the solutions for the control and worst case shock are

$$\begin{bmatrix} U_t \\ W_{t+1} \end{bmatrix} = \begin{bmatrix} -F \\ K \end{bmatrix} \begin{bmatrix} I & 0 \\ -P_{22}^{-1}P_{21} & P_{22}^{-1} \end{bmatrix} \begin{bmatrix} z_t \\ \mu_{xt} \end{bmatrix}.$$
 (16.3.17)

Using the law of motion for y_{t+1} together with (16.3.16) and (16.3.14) allows us to represent our solution recursively as

$$\begin{bmatrix} z_{t+1} \\ \mu_{x,t+1} \end{bmatrix} = \mathcal{M} \begin{bmatrix} z_t \\ \mu_{xt} \end{bmatrix}$$
(16.3.18*a*)

$$x_t = \begin{bmatrix} -P_{22}^{-1}P_{21} & P_{22}^{-1} \end{bmatrix} \begin{bmatrix} z_t \\ \mu_{xt} \end{bmatrix}.$$
 (16.3.18b)

where

$$\mathcal{M} = \begin{bmatrix} I & 0\\ P_{21} & P_{22} \end{bmatrix} (A - BF + CK) \begin{bmatrix} I & 0\\ -P_{22}^{-1}P_{21} & P_{22}^{-1} \end{bmatrix}.$$
 (16.3.19)

In summary, we solve the robust Stackelberg problem by formulating a particular optimal linear regulator, solving the associated matrix Riccati equation (16.3.4) for P, computing F, K, and then partitioning P to obtain representation (16.3.18a).

16.4. Incorporating robustness for the followers

So far we have concentrated on getting a robust rule for the leader, taking as given the Euler equations that characterize a follower's behavior. In this section, we point out that by including Euler equations for the followers's worst case shocks among the implementability constraints, we can impute a concern for robustness to the followers as well as to the leader. For a representative follower example, we shall index a follower's concern about robustness by a multiplier θ that need not equal the robustness parameter Θ of the leader.

16.4.1. An approach enabled by the Bellman-Isaacs condition

To formulate the implementability constraints concisely, we rely on findings about the zero-sum two-player dynamic game that underlies the single-agent robust control problem. Başar and Bernhard (1995) and chapter 6 show that equilibrium outcomes are identical for several games with different timing protocols for the maximizing and minimizing players. Among these different timing protocols is one in which the minimizing and maximizing players simultaneously choose entire sequences of state-contingent decisions at time 0. By using that timing protocol for a *follower's* two-person zero-sum game, we can represent a follower's decisions by the stabilizing solution of the Euler equations for extremizing his objective with respect to *both* his 'natural control' u_t and his pseudo-control w_{t+1} , the worst case shocks.¹⁰ Then in the robust Stackelberg problem, we can regard the first-order conditions of a follower, *including* his worst-case shock process, as among the implementability conditions for the Stackelberg leader. This leads to an equilibrium of the game between the leader and the follower in which each understands the decision rules of the other, and in which the leader takes into account how the follower's decisions respond to its own. To know how the follower responds, the leader has to keep track of how the worst case shocks of the follower respond to the leader's decisions. This impels us to include the worst case shock process of a follower in the state vector for the leader. The leader manipulates a follower's worst case shock process in order to manipulate his robust best response to the leader's own robust decisions.

 $^{10\,}$ A Bellman-Isaacs condition on the value function described in chapter 6 allows us to characterize the solution of the robust control problem in this way. We use a substantial and very important result about zero-sum games here. For general two-person games, the Markov-perfect equilibrium cannot be computed by stacking and solving the Euler equations for the two players. Doing that would produce a candidate equilibrium that would not be subgame perfect. But the Bellman-Isaacs condition that pertains to two-player zero-sum games implies that a Markov perfect equilibrium can be computed by stacking and solving the Euler equations. For proofs, see Başar and Bernhard (1995) and chapter 6. Technically, the irrelevance of timing protocols for zero-sum two-player dynamic games is related to Chari, Kehoe, and Prescott's (1989, pp. 269–272) that time-inconsistency in macroeconomics occurs only in situations in which there is *conflict* between a society's objective and those of the agents within it. Chari, Kehoe, and Prescott show that without such conflict, the existence of a single value function makes irrelevant the order of maximization. Comparing their result to the one we are using that is based on the Bellman-Isaacs condition for two-player zerosum dynamic games, it can be seen that to avoid time inconsistency really only requires that objective functions of different decision makers be completely *aligned*, a condition that holds when there is perfect conflict.

16.4.2. Heterogeneity of worst case models

By following this recipe, we can construct an equilibrium in which leaders and followers share a common approximating model. However, differences in their preferences will lead them to slant their worst case models in different directions away from their common approximating model, as the two types of agents use their own worst-case analyses to investigate the fragility of alternative rules to possible misspecifications of that common approximating model. In section 16.5, we illustrate our equilibrium concept with an example.

16.4.3. Evolution under approximating model and robust rules

The general equilibrium models described in chapter 10 contained a vector of exogenous state variables whose evolution was given by $e_{t+1} = A_{22}e_t + C_2\epsilon_{t+1}$.¹¹ As we shall illustrate in section 16.5, when we want to model concerns about robustness for both the leader and the follower, we shall include subvectors designed to distinguish between a leader's and a follower's worst-case views of future values of the exogenous state variables. This means that z_t includes two versions of e_t , one for the leader and one for the follower, which under the distorted models take the forms

$$e_{\ell,t+1} = A_{22}e_{\ell,t} + C_2(\epsilon_{t+1} + W_{t+1})$$

$$e_{f,t+1} = A_{22}e_{f,t} + C_2(\epsilon_{t+1} + w_{t+1})$$
(16.4.1)

where e_f pertains to the follower and e_{ℓ} pertains to the leader, W_{t+1} is the leader's distortion to the mean of the innovation, and w_{t+1} is the follower's. The law of motion (16.3.18) then contains descriptions of the evolution of both e_{ℓ} and e_f . We impose the initial condition $e_{\ell,0} = e_{f,0} = e_0$ to reflect that the Stackelberg problem conditions on time 0 information that is available to both the leader and the followers.

The solution of the Stackelberg problem (16.3.18) then gives sequences $\{z_t, \mu_{xt}, t \ge 0\}$ that respect the initial condition for z_0, μ_{x0} , where the initial condition for z_0 includes $e_{\ell,0} = e_{f,0} = e_0$. The components of the solution $\{e_{\ell,t}, e_{f,t}\}$ for $t \ge 1$ are the worst case forecasts of future e_s made by the leader and the followers, respectively, conditional on information at time 0. A

¹¹ In chapter 10, we called it z_t rather than e_t . We change notation because in this chapter we have used z_t for something else.

certainty equivalence argument can be used to show that for a stochastic version of the model driven by Gaussian shocks, these forecasts equal the conditional means of these agents' worst case models conditioned on time 0 information, including e_0 .

16.4.4. Evolution of the economy under Stackelberg solution

To understand how we can repeatedly solve the Stackelberg problem to trace out the evolution of the economy over time under commitment, it is sufficient to consider the situation at time 1 and to consider the time 1 values $z_1, \mu_{x,1}$ that were determined by the time 0 solution. It is appropriate to take the variables in $z_1, \mu_{x,1}$ other than $e_{\ell,1}, e_{f,1}$ as initial conditions for the time 1 version of the Stackelberg problem. Because at time 1 the actual value of e_1 will have been realized and known to both the leader and the followers, to pose the time 1 Stackelberg problem, it is appropriate to set $e_{f,1}$ and $e_{\ell,1}$ both equal to the realized value e_1 . This step expresses the idea that both the leader and the follower observe e_1 at time 1.¹² We then can solve a time 1 Stackelberg problem with these initial conditions, which is accomplished by simply applying solution (16.3.18) to these initial conditions.

That the Stackelberg leader is committed to the time 0 plan is expressed in solving the time 1 taking $\mu_{x,1}$ determined in the time 0 solution as initial conditions for the time 1 problem. The part of the solution for $\{e_{\ell,t}, e_{f,t}, t \ge 1\}$ corresponds to what in a stochastic version of the model will be conditional expectations of future e's from the the leader's and followers' worst case models, conditioned on time 1 information that includes e_1 .

To simulate the evolution of the state variables under the approximating model and the robust decision rules, we can dispense with e_f and e_ℓ as separate state variables and create a collapsed state called \tilde{z}_t that imposes $e_{\ell,t} = e_{f,t} = e_t$. The law of motion of this reduced state vector can be deduced from (16.3.18) and takes the form

$$\begin{bmatrix} \tilde{z}_{t+1} \\ \mu_{x,t+1} \end{bmatrix} = \mathcal{N} \begin{bmatrix} \tilde{z}_t \\ \mu_{xt} \end{bmatrix} + \begin{bmatrix} \tilde{I} & 0 \\ \tilde{P}_{21} & P_{22} \end{bmatrix} C \epsilon_{t+1}$$
(16.4.2*a*)

¹² This mechanical process of resetting the initial conditions for e_f, e_ℓ as time passes thus captures how uncertainty is resolved as time passes.

$$x_t = \begin{bmatrix} -P_{22}^{-1} \tilde{P}_{21} & P_{22}^{-1} \end{bmatrix} \begin{bmatrix} \tilde{z}_t \\ \mu_{xt} \end{bmatrix}.$$
 (16.4.2b)

where the $\tilde{\cdot}$ above matrices refer to objects whose size has been reduced by appropriately combining the coefficients of e_{ℓ} and e_f and where \mathcal{N} is obtained as follows from \mathcal{M} . First, to get the rows in \mathcal{N} pertaining to variables other than e_{ℓ} and e_f , add the coefficients on e_{ℓ} and e_f in \mathcal{M} and use them as the coefficients on e in \mathcal{N} . Second, delete the rows in \mathcal{M} corresponding to e_{ℓ} and e_f and replace them with the rows corresponding to

$$e_{t+1} = A_{22}e_t + C_2\epsilon_{t+1}.$$

The matrix P_{21} is formed from P_{21} by a similar deletion procedure.

In section 16.5.4, we shall illustrate within a concrete example this procedure for moving from representation (16.3.18) to a recursive representation of outcomes under the approximating model and the robust decision rule for the Stackelberg leader.

16.5. A monopolist with a competitive fringe

As an example, this section studies an industry with a large firm that acts as a Stackelberg leader with respect to a competitive fringe. The industry produces a single nonstorable homogeneous good. One large firm called the monopolist produces Q_t and a representative firm in a competitive fringe produces q_t . We use q_t to denote the quantity chosen by the individual competitive firm and \overline{q}_t to denote the equilibrium quantity. In equilibrium, $q_t = \overline{q}_t$, but in posing the optimum problem of the representative competitive firm in the competitive fringe takes Q_t and \overline{q}_t and \overline{q}_t . The representative firm in the competitive fringe takes Q_t and \overline{q}_t as exogenous and chooses sequentially. In light of the responses of the representative firm in the competitive fringe, the monopolist commits to a policy at time 0, taking into account its ability to manipulate the price sequence and the worst case beliefs of the representative competitive fringe's best response, the monopolist views itself as choosing \overline{q}_{t+1} and Q_{t+1} for $t \ge 0$, as well as the representative competitive firm's worst-case shock process w_{t+1} for $t \ge 0$.

Costs of production are $C_t = eQ_t + .5gQ_t^2 + .5c(Q_{t+1} - Q_t)^2$ for the monopolist and $\sigma_t = dq_t + .5hq_t^2 + .5c(q_{t+1} - q_t)^2$ for the representative competitive

firm, where d > 0, e > 0, c > 0, g > 0, h > 0 are cost parameters. There is a linear inverse demand curve

$$p_t = A_0 - A_1(Q_t + \overline{q}_t) + v_t, \qquad (16.5.1)$$

where A_0, A_1 are both positive and v_t is a disturbance to demand governed by

$$v_{t+1} = \rho v_t + C_{\epsilon} \check{\epsilon}_{t+1} \tag{16.5.2}$$

and where $|\rho| < 1$ and $\check{\epsilon}_{t+1}$ is an i.i.d. sequence of random variables with mean zero and variance 1. The monopolist and the representative competitive firm share equation (16.5.2) as their approximating model for the demand shock. The monopolist and the representative competitive firm both want decision rules that are robust to alternative specifications of the process for the demand shock. Because the monopolist and the representative firm in the competitive fringe potentially have different worst case models of the demand shock, we distinguish between them by letting v_t denote the process perceived by the representative firm, and V_t the process perceived by the monopolist. For the representative competitive firm, the alternative models of the demand shock have the form

$$v_{t+1} = \rho v_t + C_{\epsilon} (\epsilon_{t+1} + w_{t+1}). \tag{16.5.3}$$

For the monopolist, they have the form

$$V_{t+1} = \rho V_t + C_{\epsilon} (\tilde{\epsilon}_{t+1} + W_{t+1}). \tag{16.5.4}$$

Here w_{t+1}, W_{t+1} are specification errors for the representative competitive firm and the monopolist, respectively, and $\epsilon_{t+1}, \tilde{\epsilon}_{t+1}$ are other i.i.d. random processes with mean zero and variance 1. The distortion W_{t+1} can feed back on the history of the state of the market, namely, (\bar{q}, Q, v, V) and the distortion w_{t+1} can feed back on the history of the market-wide version of the state (\bar{q}, Q, v, V) .¹³ The distortions w_{t+1} and W_{t+1} will typically differ because the monopolist and the representative competitive firm have different objectives.

¹³ In section 16.5.3, we show that in a recursive representation of the problem of a representative firm in the competitive fringe, the household's decision rule feeds back on a 'big K' version of the state (\bar{q}, Q, v, V) and also the individual firm's quantity q and v.

16.5.1. The competitive fringe

The representative competitive firm regards $\{Q_t, \overline{q}_t\}_{t=0}^{\infty}$ as given stochastic processes and chooses an output plan $\{q_{t+1}\}_{t=0}^{\infty}$ and shock distortion process $\{w_{t+1}\}_{t=0}^{\infty}$ to extremize

$$E_0 \sum_{t=0}^{\infty} \beta^t \left\{ p_t q_t - \sigma_t + \beta \theta w_{t+1}^2 \right\}, \quad \beta \in (0, 1)$$
 (16.5.5)

subject to q_0 given, where E_t is the mathematical expectation based on time tinformation evaluated with respect to a distorted model that includes (16.5.3). Here θ is the robustness parameter of the representative firm in the competitive fringe, which can differ from Θ , the robustness parameter of the monopolist. Let $u_t = q_{t+1} - q_t$. We take (u_t, w_{t+1}) as the representative competitive firm's composite control vector at t. Subject to (16.5.1) and (16.5.3), first orderconditions for extremizing (16.5.5) with respect to u_t, w_{t+1} , respectively, are

$$u_{t} = E_{t}\beta u_{t+1} - c^{-1}\beta hq_{t+1} + c^{-1}\beta E_{t}(p_{t+1} - d)$$

$$w_{t+1} = -\frac{1}{2\theta}C_{\epsilon}E_{t}q_{t+1} + \beta\rho E_{t}w_{t+2}$$
(16.5.6)

for $t \geq 0$.

In more detail, we derive the first-order conditions (16.5.6) by forming the following Lagrangian for the representative firm in the competitive fringe:

$$L = E_0 \sum_{t=0}^{\infty} \beta^t \Big\{ [A_0 - A_1(Q_t + \overline{q}_t) + v_t] q_t - [dq_t + .5hq_t^2 + .5cu_t^2] \\ \beta \theta w_{t+1}^2 + \ell_{1t} [q_t + u_t - q_{t+1}] + \ell_{2t} [\rho v_t + C_\epsilon w_{t+1} - v_{t+1}] \Big\}.$$
(16.5.7)

Here $\{\ell_{1t}, \ell_{2t}\}$ are sequences of Lagrange multipliers. Taking $\{Q_t, \overline{q}_t\}_{t=0}^{\infty}$ as given, the representative firm maximizes L with respect $\{u_t, q_{t+1}\}_{t=0}^{\infty}$ and minimizes it with respect to $\{w_{t+1}, v_{t+1}\}_{t=0}^{\infty}$. Rearranging the first order conditions for (u_t, q_{t+1}) gives the first equation of (16.5.6), while rearranging the first-order conditions for (w_{t+1}, v_{t+1}) gives the second equation of (16.5.6), which from now on we call the Euler equation for w_{t+1} .

We can appeal to a certainty equivalence principle to justify working with a non-stochastic version of (16.5.6) that we form by dropping the expectation operator and the random terms $\check{\epsilon}_{t+1}$ and ϵ_{t+1} from (16.5.2) and (16.5.3).¹⁴ Shift (16.5.1) forward one period, set $q_t = \overline{q}_t$ for all $t \ge 0$, and substitute for p_{t+1} in (16.5.6) to get

$$u_{t} = \beta u_{t+1} - c^{-1}\beta h \overline{q}_{t+1} + c^{-1}\beta (A_{0} - d) - c^{-1}\beta A_{1} \overline{q}_{t+1} - c^{-1}\beta A_{1} Q_{t+1} + c^{-1}\beta v_{t+1}$$

$$w_{t+1} = -\frac{1}{2\theta} C_{\epsilon} \overline{q}_{t+1} + \beta \rho w_{t+2}.$$
(16.5.8)

Equation (16.5.8) combines the Euler equations of the representative firm in the competitive fringe with market clearing.¹⁵ Note that v, and not V, appears in the first equation of (16.5.8). This reflects how the representative competitive firm's forecasts influence its decisions, a fact that the monopolist will acknowledge when he designs his policy.

¹⁴ Here we use a method that Sargent (1987) used to compute a rational expectations equilibrium. The key step is to eliminate price and output by setting $q_t = \overline{q}_t$ and substituting from the inverse demand curve and the production function into the firm's first-order conditions to get a difference equation in capital.

¹⁵ As shown in Sargent (1987) in the case without robustness, (16.5.8) is also the Euler equation for a fictitious planner who takes Q_t as exogenous and who chooses a sequence for $\{q_{t+1}\}_{t=0}^{\infty}$ to maximize the discounted sum of consumer and producer surplus. Given stable sequences $\{Q_t, v_t\}$, we could solve (16.5.8) and $u_t = \overline{q}_{t+1} - \overline{q}_t$ to express the competitive fringe's output sequence as a function of the monopolist's output sequence.

16.5.2. The monopolist's problem

The monopolist views the sequence of Euler equations-cum-market-clearing conditions (16.5.8) as implementability constraints. We can represent the constraints impinging on the monopolist, including (16.5.8), in terms of the transition law:

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -\frac{1}{2\theta}C_{\epsilon} & 0 & \beta\rho \end{bmatrix} \begin{bmatrix} 1 \\ v_{t+1} \\ Q_{t+1} \\ \overline{q}_{t+1} \\ u_{t+1} \\ w_{t+2} \end{bmatrix}$$

$$= \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & \rho & 0 & 0 & 0 & 0 \\ 0 & \rho & 0 & 0 & 0 & 0 \\ 0 & \rho & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ v_t \\ Q_t \\ \overline{q}_t \\ u_t \\ w_{t+1} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} U_t + \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} W_{t+1},$$
(16.5.9)

where $U_t = Q_{t+1} - Q_t$ is the control of the monopolist. The last row portrays (16.5.8). Represent (16.5.9) as

$$y_{t+1} = Ay_t + BU_t + CW_{t+1}.$$
(16.5.10)

Although we have included (u_t, w_{t+1}) as components of the 'state' y_t in the monopolist's transition law (16.5.10), (u_t, w_{t+1}) are actually 'jump' variables that correspond to x_t in section 16.3. The analysis in section 16.3 implies that the solution of the monopolist's problem is encoded in the Riccati equation associated with a robust linear regulator that takes (16.5.10) as the transition law.

To match the setup of section 16.3, we partition y_t as $y'_t = \begin{bmatrix} z'_t & x'_t \end{bmatrix}$ where $z'_t = \begin{bmatrix} 1 & v_t & V_t & Q_t & \overline{q}_t \end{bmatrix}$, $x'_t = \begin{bmatrix} u'_t & w'_{t+1} \end{bmatrix}$, and let $\mu'_{xt} = \begin{bmatrix} \mu_{ut} & \mu_{wt} \end{bmatrix}$ be the vector of multipliers associated with the Euler equations for (u_t, w_{t+1}) . The monopolist's artificial optimal linear regulator problem can be expressed

$$\max_{\{U_t\}} \min_{\{W_{t+1}\}} \sum_{t=0}^{\infty} \beta^t \left\{ p_t Q_t - \mathcal{C}_t + \beta \Theta W'_{t+1} W_{t+1} \right\}$$
or

$$\max_{\{U_t\}} \min_{\{W_{t+1}\}} \sum_{t=0}^{\infty} \beta^t \left\{ (A_0 - A_1(\overline{q}_t + Q_t) + V_t) Q_t - eQ_t - .5gQ_t^2 - .5cU_t^2 + \beta \Theta W_{t+1}^2 \right\}$$
(16.5.11)

subject to (16.5.10). Notice how the monopolist's perceived demand shock V_t appears in (16.5.11). The monopolist's problem can be written

$$\max_{\{U_t\}} \min_{\{W_{t+1}\}} - \sum_{t=0}^{\infty} \beta^t \left\{ y_t' Q y_t + U_t' R U_t - \beta \Theta W_{t+1}^2 \right\}$$
(16.5.12)

subject to (16.5.10) where

and $R = \frac{c}{2}$. The results of section 16.3 apply.

16.5.3. Recursive representation of a competitive firm's problem

By using a counterpart to the 'Big K, little k' trick used in chapter 11 on page 283. The representative competitive firm takes Q, \overline{q} as exogenous processes that are determined by the following 'big K' counterpart to (16.3.18*a*):

$$Q_t = e_Q X_t$$

$$\bar{q}_t = e_q X_t$$

$$X_{t+1} = \mathcal{M} X_t$$

(16.5.13)

where

$$X_t = \begin{bmatrix} Z_t \\ M_{x,t} \end{bmatrix}$$
(16.5.14)

In addition to (16.5.13), (16.5.14), the representative firm in the competitive fringe faces the laws of motion

$$q_{t+1} = q_t + u_t$$

$$v_{t+1} = \rho v_t + C_e w_{t+1}$$
(16.5.15)

and seeks to choose a $\{w_{t+1}, v_{t+1}\}$ process to minimize and a $\{u_t, q_{t+1}\}$ sequence to maximize

$$\sum_{t=0}^{\infty} \beta^t \Big\{ \Big[A_0 - A_1 (Q_t + \bar{q}_t) + v_t \Big] q_t - \sigma_t + \beta \theta w_{t+1}^2 \Big\}$$
(16.5.16)

subject to (16.5.13), (16.5.15). This problem can be formulated as a discounted robust linear regulator with state $\tilde{X}_t = \begin{bmatrix} X'_t & q'_t & v'_t \end{bmatrix}'$ and extremizing decision rules

$$u_t = d_u \tilde{X}_t \tag{16.5.17}$$

$$w_{t+1} = d_w \tilde{X}_t. \tag{16.5.18}$$

If we add the coefficients on the two components corresponding to the state v in \tilde{X} , and also add the coefficients corresponding to the state q and \bar{q} , then it should be true that the decision rules (16.5.17), (16.5.18) for u, w that depend on X_t equal those given by (16.3.18b). In the Matlab program compfringe.m that will be used to compute the example in section 16.5.5, we verify that the decision rules line up in this way.

16.5.4. Representation of the monopolist's decision rule under the approximating model

The monopolist's decision rule has representation (16.3.18*a*), (16.3.18*b*) where $z = \begin{bmatrix} 1 & v' & V' & Q & \bar{q} \end{bmatrix}'$ and $\mu_x = \begin{bmatrix} \mu_u & \mu_w \end{bmatrix}'$. As anticipated in section 16.4.3, we have included within the state z two versions of the demand shock v, V for the monopolist, and v for the representative firm in the competitive fringe, because distinguishing between them allows us to represent the different ways that the two types of firms distort the conditional distributions of future values of the demand shock. As we have seen, the differing motives of the two types of firms inspire their respective evil agents to twist their worst case distributions of future v's in different directions.

However, because it can be constructed from observed outputs and the price, v is actually an observable state variable. As explained in section 16.4.3, we must recognize this when we simulate the approximating model under the robust decision rules. We accomplish this by reducing the description of the state

under the approximating model to $\tilde{x} = \begin{bmatrix} 1 & v & Q & \bar{q} & \mu_u & \mu_w \end{bmatrix}'$ and writing

$$\tilde{x}_{t+1} = \mathcal{N}\tilde{x}_t + \tilde{C}\epsilon_{t+1} \tag{16.5.19}$$

where the second row of (\mathcal{N}, \tilde{C}) expresses the law of motion $v_{t+1} = \rho v_t + C_e \epsilon_{t+1}$. Evidently, we can form the remaining rows of \mathcal{N} from the corresponding rows of \mathcal{M} by adding the column elements corresponding to the coefficients on v and V in \mathcal{M} to form the coefficient on v in \mathcal{N} , leaving the coefficients on the other variables the same; the remaining rows of \tilde{C} equal zero.

16.5.5. Numerical example

This section briefly describes a numerical example of the monopoly-competitive fringe model in which we start without concerns about robustness, then study the effects of successively turning on concerns about robustness for one type of agent, but not the other, and then turning them on for both.¹⁶ For parameter settings $(A_0, A_1, \rho, C_{\epsilon}, c, d, e, g, h, \beta) = (100, 1, .8, .2, 1, 20, 20, 1, 1, .95)$, Table 16.5.1 displays steady state values associated with four pairs of settings for (Θ, θ) under the approximating model and the robust rule. To represent little or no concern about robustness, we set θ or Θ equal to 100000. To activate concerns about robustness, we set θ or Θ equal to 10.

The first column of Table 16.5.1 serves as a benchmark, concerns about robustness having been turned off for both the monopolist and the competitive firms by setting $\Theta = \theta \approx +\infty$. The next two columns turn on a concern about robustness for one but not the other of the two types of agents, while the fourth column turns on a concern about robustness for both types. The entries in the table show that a main effects of activating concerns about robustness are to make the steady state values of the worst case shocks w and W negative. In effect, firms' pessimistic forecasts about demand push their outputs down. In the middle two columns in the table in which concerns about robustness are turned on for one but not the other type of firm, the type with the concern about robustness produces less and the other type produces more than under the benchmark steady state without concerns about robustness.

 $^{^{16}}$ The calculations are performed by the Matlab programs compfringall.m and compfringefn.m.

(Θ, θ)	(∞,∞)	$(\infty, 10)$	$(10,\infty)$	(10, 10)
p	50	50.23	50.07	50.03
q	30	29.72	30.10	29.82
Q	20	20.05	19.83	19.88
w	0	247	0	25
W	0	0	0.82	83

 Table 16.5.1:
 Steady state values

16.6. Concluding remarks

This chapter has generalized standard methods for solving Ramsey problems in linear-quadratic forward looking models to include a common concern for model misspecification to both the government and private agents. The government and private agents share an approximating model that describes the shocks and other exogenous variables hitting the economy. We add two parameters Θ and θ to the standard rational expectations setup, penalty parameters that measure sets of models near the approximating model over which the leader and the followers, respectively, want robust decision rules. We compute the Ramsey rule by forming an optimal linear regulator problem and carefully exchanging the roles of the forward looking model's artificial state variables and the Lagrange multipliers on their laws of motion. Mechanically, robustness for the leader is achieved simply by adding another control to the regulator problem, a distortion to the conditional mean of the disturbances that is chosen by a fictitious evil agent.

A. Invariant subpace method

Let $L = L^* \beta^{-.5}$ and transform the system (16.3.12) to

$$L\begin{bmatrix} y_{t+1}^*\\ \mu_{t+1}^* \end{bmatrix} = N\begin{bmatrix} y_t^*\\ \mu_t^* \end{bmatrix},$$
(16.A.1)

where $y_t^* = \beta^{t/2} y_t, \mu_t^* = \mu_t \beta^{t/2}$. Now $\lambda L - N$ is a symplectic pencil, so that the generalized eigenvalues of L, N occur in reciprocal pairs: if λ_i is an eigenvalue, then so is λ_i^{-1} .

We can use Evan Anderson's Matlab program schurg.m to find a stabilizing solution of system (16.A.1). The program computes the ordered real generalized Schur decomposition of the matrix pencil. Thus, schurg.m computes matrices \bar{L}, \bar{N}, V such that \bar{L} is upper triangular, \bar{N} is upper block triangular, and V is the matrix of right Schur vectors such that for some orthogonal matrix W the following hold:

$$WLV = L$$

$$WNV = \bar{N}.$$
(16.A.2)

Let the stable eigenvalues (those less than 1) appear first. Then the stabilizing solution is

$$\mu_t^* = P y_t^* \tag{16.A.3}$$

where

$$P = V_{21} V_{11}^{-1},$$

 V_{21} is the lower left block of V, and V_{11} is the upper left block.

If L is nonsingular, we can represent the solution of the system as 17

$$\begin{bmatrix} y_{t+1}^* \\ \mu_{t+1}^* \end{bmatrix} = L^{-1} N \begin{bmatrix} I \\ P \end{bmatrix} y_t^*.$$
(16.A.4)

 $17\,$ The solution method in the text assumes that L is nonsingular and well conditioned. If it is not, the following method proposed by Evan Anderson will work. We want to solve for a solution of the form

$$y_{t+1}^* = A_o^* y_t^*$$

Note that with (16.A.3),

$$L[I;P]y_{t+1}^* = N[I;P]y_t^*$$

The solution A_o^* will then satisfy

$$L[I;P]A_o^* = N[I;P].$$

Thus A^{o*} can be computed via the Matlab command

$$A_o^* = (L * [I; P]) \setminus (N * [I; P]).$$

The solution is to be initiated from (16.A.3). We can use the first half and then the second half of the rows of this representation to deduce the following recursive solutions for y_{t+1}^* and μ_{t+1}^* :

$$y_{t+1}^* = A_o^* y_t^*$$

$$\mu_{t+1}^* = \psi^* y_t^*$$
(16.A.5)

Now express this solution in terms of the original variables:

$$y_{t+1} = A_o y_t$$

 $\mu_{t+1} = \psi y_t,$
(16.A.6)

where $A_o = A_o^* \beta^{-.5}, \psi = \psi^* \beta^{-.5}$. We also have the representation

$$\mu_t = P y_t. \tag{16.A.7}$$

The matrix $A_o = A - \tilde{B}F$, where F is the matrix for the optimal decision rule.

B. The Riccati equation

16.B.1. The Riccati equation

The stabilizing P obeys a Riccati equation coming from the Bellman equation. Substituting $\mu_t = Py_t$ into (16.3.10) and (16.3.8b) gives

$$(I + \beta \tilde{B} \tilde{R}^{-1} \tilde{B} P) y_{t+1} = A y_t \tag{16.B.1a}$$

$$\beta A' P y_{t+1} = -Q y_t + P y_t.$$
 (16.B.1b)

A matrix inversion identity implies

$$(I + \beta \tilde{B} \tilde{R}^{-1} \tilde{B}' P)^{-1} = I - \beta \tilde{B} (\tilde{R} + \beta \tilde{B}' P \tilde{B})^{-1} \tilde{B}' P.$$
(16.B.2)

Solving (16.B.1a) for y_{t+1} gives

$$y_{t+1} = (A - \tilde{B}F)y_t \tag{16.B.3}$$

where

$$F = \beta (\tilde{R} + \beta \tilde{B}' P \tilde{B})^{-1} \tilde{B}' P A.$$
(16.B.4)

Pre multiplying (16.B.3) by $\beta A'P$ gives

$$\beta A' P y_{t+1} = \beta (A' P A - A' P \tilde{B} F) y_t. \tag{16.B.5}$$

For the right side of (16.B.5) to agree with the right side of (16.B.1b) for any initial value of y_0 requires that

$$P = Q + \beta A' P A - \beta^2 A' P \tilde{B} (\tilde{R} + \beta \tilde{B}' P \tilde{B})^{-1} \tilde{B}' P A.$$
(16.B.6)

Equation (16.B.6) is the algebraic matrix Riccati equation associated with the ordinary linear regulator for the system $A, \tilde{B}, Q, \tilde{R}$.

C. Another Bellman equation

We briefly indicate the connection of the preceding formulation to that of Kydland and Prescott (1980) and Marcet and Marimon (2000). For a class of problems with structures close to ours, they construct a Bellman equation in a state vector defined as (z, μ_x) : these are the 'natural' state variables and the vector of multipliers on the laws of motion for the 'jump' variables x_t . We show how to modify that Bellman equation to include a concern about model misspecification.

Let μ_{xt} denote the sub vector of multipliers attaching to the implementability constraints that summarize the Euler equations of the private sector. Then the Lagrangian for the optimum problem (16.3.7) can be written

$$\mathcal{L} = -\sum_{t=0}^{\infty} \beta^{t} \left\{ \begin{bmatrix} z_{t} \\ x_{t} \end{bmatrix}' Q \begin{bmatrix} z_{t} \\ x_{t} \end{bmatrix} + U_{t}' R U_{t} - \beta \theta w_{t+1}' w_{t+1} + \beta \mu_{x,t+1}' (A_{21}z_{t} + A_{22}x_{t} + B_{2}U_{t} + C_{2}w_{t+1} - x_{t+1}) \right\}.$$
(16.C.1)

This Lagrangian is to be 'extremized' (i.e., maximized or minimized, as appropriate) with respect to sequences $\{z_t, x_t \mu_{x,t}, w_{t+1}\}$ subject to $\lambda_0 = 0$ and the transition law

$$z_{t+1} = A_{11}z_t + A_{12}x_t + B_1U_t + C_1w_{t+1}.$$
(16.C.2)

Equation (16.C.1) can be rewritten

$$\mathcal{L} = -\sum_{t=0}^{\infty} \beta^{t} \left\{ \begin{bmatrix} z_{t} \\ x_{t} \end{bmatrix}' Q \begin{bmatrix} z_{t} \\ x_{t} \end{bmatrix} + U_{t}' R U_{t} - \beta \theta w_{t+1}' w_{t+1} + (\beta \mu_{x,t+1}' A_{22} - \mu_{x,t}') x_{t} + \beta \mu_{x,t+1}' (A_{21} z_{t} + B_{2} U_{t} + C_{2} w_{t+1}) \right\},$$
(16.C.3)

which is to be extremized with respect to the same constraints (16.C.2). Define the one-period return function $-\tilde{r}(z,\mu_x,x,\mu_x^*,w) = \begin{bmatrix} z \\ x \end{bmatrix}' Q \begin{bmatrix} z \\ x \end{bmatrix} + u'Ru - \theta w'w + b'w'w'$

 $(\beta \mu_x^{*'} A_{22} - \mu_x')x + \beta \mu_x^{*'} (A_{21}z + B_2u + C_2w)$, where * superscripts denote one-period ahead values. Let $v(z, \mu_x)$ be the optimum value of the problem starting with augmented state (z, μ_x) . Problem (16.C.3) is recursive and has the following Bellman equation:

$$v(z,\mu_x) = \max_{\{u,x\}} \min_{\{w,\mu_x^*\}} \left\{ \tilde{r}(z,\mu_x,x,\mu_x^*,w) + \beta v(z^*,\mu_x^*) \right\}$$
(16.C.4)

where the extremization is subject to

$$z^* = A_{11}z + A_{12}x + B_1u + C_1w. (16.C.5)$$

The Bellman equation (16.C.4), (16.C.5) is a version of the recursive saddle problem described by Kydland and Prescott (1980) and Marcet and Marimon (2000). We have added a concern for robustness via the extra minimization with respect to the shock distortion w. In related contexts, Marcet and Marimon stress that while such problems are not recursive in the natural state variables z alone, they becomes recursive when the multipliers μ_x are included.

Although one could solve our problem by iterating to convergence on (16.C.4), (16.C.5), it is more convenient for us to use the method described in section 16.3 that solves the Riccati equation (16.3.4) and its associated Bellman equation.

Chapter 17 Non-linear models

17.1. Introduction

This chapter discusses a preference for robustness in settings that extend beyond the linear-quadratic examples we have concentrated on up to now. We permit both the return function and the transition law to be of general functional forms. The state vector obeys a Markov process. A preference for robustness to model misspecification can be expressed by altering the conditional expectation operator in the Bellman equation. That new operator is connected with the risk-sensitive control theory of chapter 6.

17.1.1. The \mathcal{R} operator for LQ problems

Chapters 6 and 7 showed how for a linear quadratic problem, a robust decision rule can be found by iterating to convergence on a composite operator $T \circ \mathcal{D}$ in place of the ordinary operator T defined by the right side of the Riccati equation for the matrix P in the value function -x'Px - p. Thus, for the nonstochastic linear-quadratic case studied in chapter 6, a Bellman equation that induces a robust rule is

$$-x'Px = \max_{u} \left[r(x, u) - \beta y' \mathcal{D}(P)y \right]$$
(17.1.1)

where r(x, u) = -x'Qx - u'Ru, the maximization is subject to $y = A_o x + Bu$, and

$$\mathcal{D}(P) = P + \theta^{-1} P C (I - \theta^{-1} C' P C)^{-1} C' P.$$

The \mathcal{D} operator in Bellman equation (17.1.1) rewards robust decision rules. The operator \mathcal{D} verifies the following equality:

$$J \equiv -x'A'\mathcal{D}(P)Ax = \min_{w} \left[\theta w'w - (Ax + Cw)'P(Ax + Cw)\right].$$
 (17.1.2)

The problem on the right is to minimize $\theta w'w + y'Py$ subject to the approximating model y = Ax + Cw, where y'Py is the continuation value function

-393 -

of next period's state y. Thus, in iterating to convergence on $T \circ \mathcal{D}$, the \mathcal{D} operator reflects the actions of the minimizing agent who distorts models, and the T operator reflects the actions of the maximizing agent.

Chapter 7 reinterprets the operator \mathcal{D} in terms of risk-sensitivity where the law of motion is $y = A_o x + Bu + C\epsilon$ and ϵ is a Gaussian vector with mean zero and identity covariance matrix. Define $\alpha = 2\theta$.¹ The Bellman equation for the value function V(x) for a risk-sensitive control problem is

$$V(x) = \max_{u} \{ r(x, u) + \beta \mathcal{R}(V(y))(x) \}$$
(17.1.3)

where

$$\mathcal{R}(V(y))(x) = -\alpha \log E \exp\left(\frac{-V(y)}{\alpha}\right) \left| x.$$
(17.1.4)

For the fixed control law u = -Fx and the quadratic continuation value function V(y) = -y'Py - p, we have

$$\mathcal{R}(V(y))(x) = -x'A'\mathcal{D}(P)Ax - p - \theta^{-1}\log\det(I - \theta^{-1}C'PC).$$

The \mathcal{D} on the right side encapsulates a 'twisting' of probabilities induced by risk-sensitive preferences.

In the linear-quadratic case studied in chapters 6 and 7, the robust decision rule associated with a given $\theta > 0$ attains the value function V(x) that solves the risk-sensitive Bellman equation (17.1.3). Thus robustness to model misspecification can be achieved by replacing the expectations operator E with the distorted expectations operator \mathcal{R} on the right side of the usual Bellman equation, as in (17.1.3).

This chapter shows how these ideas extend beyond linear-quadratic problems. By using \mathcal{R} as defined in the first line of (17.1.4) to replace the conditional expectations operator in the Bellman equation, decision rules that are robust to model misspecification can be computed easily.

¹ In earlier chapters, we typically used θ as the multiplier on w'w, where w'w equaled two times entropy. Here we use α as the corresponding multiplier on entropy.

17.1.2. Markov perturbations

In this chapter, an approximating model is a controlled Markov chain. To express a concern about model misspecification, we suppose that a decision maker believes that an unknown member of a set of unspecified nearby models generates the data. We form the set of models by perturbing the Markov transition density of an approximating model. This way of proceeding lets us avail ourselves of formalizations for Markov processes from large deviation theory (*e.g.* see Dupuis and Ellis, 1997).

The decision maker's approximating model is a Markov process with state $x \in X$ and time-invariant transition density $\pi(x', x)$, where x denotes the state today and x' the state tomorrow. We form a perturbed model by multiplying π by a function w(x', x) > 0, then rescaling appropriately to make the resulting object a transition density:

$$\pi^{w}(x',x) = \frac{w(x',x)\pi(x',x)}{\int w(x',x)\pi(x',x)dx'}.$$
(17.1.5)

The likelihood ratio, or Radon-Nikodym derivative of π^w with respect to π , is evidently

$$\frac{w(x',x)}{\int w(x',x)\pi(x',x)dx'}.$$
(17.1.6)

From w(x', x) > 0, it follows that π^w puts positive probability on the same events as does π (i.e., π^w is said to be absolutely continuous with respect to π). This means that statistically π^w can be difficult to distinguish from π using a finite number of observations.

It is sometimes convenient to characterize a Markov process with a conditional expectations operator \mathcal{T} that is defined as follows. For any test function ϕ belonging to the class Φ of bounded continuous functions, define the operator \mathcal{T} by

$$\mathcal{T}(\phi)(x) = E\left[\phi(x_{t+1})|x_t = x\right].$$
(17.1.7)

For a rich enough set of functions in Φ , the conditional expectations operator \mathcal{T} characterizes the transition density π . The expectations operator associated with the distorted transition density π^w in (17.1.5) can be expressed

$$\mathcal{T}^w(\phi) = \frac{\mathcal{T}(w\phi)}{\mathcal{T}(w)}.$$

Let $E^w(\cdot|x)$ denote the mathematical expectation with respect to the distorted model.

17.1.3. Relative entropy

To embody the idea that the approximating model is good, we want a convenient way to measure discrepancy from the approximating model. We measure discrepancy by 'relative entropy,' defined as the expected value of the log-likelihood ratio conditional on x, where the conditional expectation is evaluated with respect to the density associated with the twisted density π^w .² As in chapter 2, we define relative entropy for a candidate model indexed by w as

$$I(w)(x) \equiv E^{w} \left[\log \frac{\pi^{w}(x',x)}{\pi(x',x)} \middle| x \right]$$

= $E^{w} \left[\log \frac{w(x',x)}{\mathcal{T}(w)(x)} \middle| x \right]$
= $\mathcal{T}^{w} (\log w)(x) - \log[\mathcal{T}(w)(x)] \ge 0.$ (17.1.8)

Relative entropy is not a metric because it treats the approximating model π and the alternative model π^w asymmetrically. The asymmetry emerges because the expectation is evaluated with respect to the 'twisted' distribution π^w . Relative entropy is prominent in both information theory and large deviation theory and satisfies several attractive properties:³ I(w) nonnegative, but I(w) = 0 if w is constant.⁴ Substituting for \mathcal{T}^w in (17.1.8) gives:

$$I(w) = \frac{\mathcal{T}[w\log(w)]}{\mathcal{T}(w)} - \log[\mathcal{T}(w)]$$

= $E\left[\frac{w(x',x)}{E[w(x',x)|x]}\log\left(\frac{w(x',x)}{E[w(x',x)|x]}\right)|x\right].$ (17.1.9)

² For readers of Dupuis and Ellis (1997, Chapter 1, Section 4), think of the transition density associated with \mathcal{T} as Dupuis and Ellis's θ ; and think of $\frac{w(z)}{\mathcal{T}(w)(y)}$ as Dupuis-Ellis's Radon-Nikodym derivative $\frac{d\gamma}{d\theta}$. For Dupuis and Ellis, relative entropy is $\int \log\left(\frac{d\gamma}{d\theta}\right) d\gamma$.

 $^{^3}$ See Shore and Johnson, 1980 and Csiszar, 1991 for axiomatic justifications.

⁴ For Markov specifications with stochastically singular transitions, $\frac{w(x',x)}{T(w)(x)}$ may be one even when w is not constant. For these systems, we have in effect over parameterized the perturbations, although in a harmless way.

As indicated in chapter 8, we limit concern about robustness by insisting that the robust decision maker pay special attention to models with small relative entropies because they are difficult to distinguish empirically from the approximating model.

17.2. Value function for robustness

Given a current-period reward function U(x) and a known Markov process, a value function W(x) for a discounted infinite horizon solves the functional equation

$$W(x) = U(x) + \beta \mathcal{T}(W)(x) \tag{17.2.1}$$

where $\beta \in (0, 1)$ is a discount factor. To attain a recursive representation while incorporating a concern about model misspecification, we replace the conditional expectations operator \mathcal{T} in (17.2.1) with an alternative transformation \mathcal{R} of the continuation value function. The operator \mathcal{R} distorts the conditional expectation operator \mathcal{T} with a single parameter $\alpha > 0$ and is defined as

$$\mathcal{R}(W) = -\alpha \log\left(\mathcal{T}\left[\exp\left(\frac{-W}{\alpha}\right)\right]\right). \tag{17.2.2}$$

The parameter $\alpha \equiv 2\theta$ is restricted to be nonnegative; as it diverges to ∞ , \mathcal{R} becomes the conditional expectation operator \mathcal{T} . The so-called risk sensitivity parameter is $\sigma = -2\alpha^{-1} = -\theta^{-1}$.⁵ In the absence of discounting, replacing \mathcal{T} with \mathcal{R} in (17.2.1) delivers the risk sensitive evaluation used in control theory.

To express a preference for robustness, we propose to iterate on the following recursion:

$$V(x) = U(x) + \beta \mathcal{R}(W)(x). \tag{17.2.3}$$

Here W(y) is a continuation value function and V(x) is a current value function.

To help motivate (17.2.3), we present an inequality that bounds how much the conditional expectation of a continuation value function deteriorates across

⁵ See Whittle (1990, 1996). As a formulation of recursive utility in the style of Epstein and Zin (1989), Weil (1993) used \mathcal{R} to make risk adjustments in a value function recursion that is not additively separable, in contrast to (17.2.1). For Weil's formulation, there exists a transformation of the value function that has a recursion that is additively separable, but the corresponding risk adjustment is different.

different probability specifications. Assume that $\alpha > 0$ and consider the following problem:

Problem A:

$$\inf_{w>0} J(w) \tag{17.2.4a}$$

where 6

$$J(w) \equiv \alpha I(w) + \mathcal{T}^{w}(W). \qquad (17.2.4b)$$

The first term on the right of (17.2.4b) is a weighted entropy measure and the second is the expectation of the continuation value function using the twisted probability model indexed by w. The objective is to find a worst-value model w, where the departures w from the approximating model are penalized at a utility-price α applied to their relative entropy. Increasing the absolute magnitude of α increases the penalty for deviating from the approximating model.

Theorem 17.2.1. Assume that \mathcal{T} can be evaluated at $\exp\left(-\frac{W}{\alpha}\right)$. For any constant k > 0, a solution to Problem A is:

$$w^* = k \exp\left(-\frac{W}{\alpha}\right),$$

which attains the minimized value

$$J(w^*) = \mathcal{R}(W)$$

where

$$\mathcal{R}(W) = -\alpha \log \left(\mathcal{T} \left[\exp \left(-\frac{W}{\alpha} \right) \right] \right)$$

The solution w^* is not unique (any k > 0 works), but the minimized value of the objective is unique and so is the associated probability law.

Proof. ⁷ To verify that w^* is the solution, write:

$$I(w) = I^*(w/w^*) + \frac{\mathcal{T}(w\log w^*)}{\mathcal{T}(w)} - \log \mathcal{T}(w^*)$$

where

$$I^*(w) = \frac{\mathcal{T}^*(w\log w)}{\mathcal{T}^*(w)} - \log \mathcal{T}^*(w)$$

⁶ Note how (17.2.4b) generalizes (17.1.2).

 $^{^{7}}$ This proof emulates the proof of Proposition 1.4.2 in Dupuis and Ellis (1997).

and

$$\mathcal{T}^*\phi \equiv \frac{\mathcal{T}(w^*\phi)}{\mathcal{T}(w^*)}.$$

Notice that I^* is itself interpretable as a measure of relative entropy and hence $I^*(w/w^*) \ge 0$. Thus the criterion J satisfies the inequality:

$$\begin{split} J(w) &= \alpha [I^*(w/w^*) + \frac{\mathcal{T}(w\log w^*)}{\mathcal{T}(w)} - \log \mathcal{T}(w^*)] + \mathcal{T}^w(W) \\ &\geq \alpha [\frac{\mathcal{T}(w\log w^*)}{\mathcal{T}(w)} - \log \mathcal{T}(w^*)] + \mathcal{T}^w(W) \\ &= -\alpha \log \mathcal{T}[\exp(-\frac{W}{\alpha})] \\ &= J(w^*). \end{split}$$

Equation (17.2.4b) implies an inequality in terms of robust evaluations of value functions.

Corollary 17.2.1. The conditional expectation of the value function W evaluated under \mathcal{T}^w satisfies the bound

$$\mathcal{T}^{w}(W) \ge \mathcal{R}(W) - \alpha I(w). \tag{17.2.5}$$

Proof. This follows immediately from $J(w^*) = \mathcal{R}(W)$ and the definition of J(w).

The first term on the right depends on α , but not on the alternative model parameterized by w. The second term is $-\alpha$ times entropy. Thus, inequality (17.2.5) justifies interpreting α as a type of *utility price of robustness*.

The robust value function W solves the functional equation:

$$W(x) = \inf_{w} \{ U(x) + \beta [\alpha I(w) + \mathcal{T}^{w}(W)(x)] \}.$$
 (17.2.6)

This can also be expressed as

$$W(x) = U(x) + \beta \mathcal{R}(W)(x). \tag{17.2.7}$$

These equations display how the continuation value function is adjusted for fear of possible model misspecification.

17.2.1. Gaussian example

Theorem 17.2.1 shows that the worst distorted transition measure obeys

$$\pi^{w}(x',x) \propto \pi(x',x) \exp(\frac{-W(x')}{\alpha}).$$
 (17.2.8)

To link this result to the linear-quadratic-Gaussian setting of earlier chapters, assume that the continuation value function is $W(x) = -x'Px - \rho$ and that $\pi(x', x)$ is Gaussian with mean A^*x and conditional covariance matrix C'C, so that x' can be represented as $x' = \mu + C\epsilon$ where $\mu = A^*x$ and ϵ is a Gaussian random vector with mean zero and identity covariance matrix. Using the definition $\sigma = -\theta^{-1} = -2\alpha^{-1}$ and the preceding assumption about the conditional distribution of x', (17.2.8) implies

$$\pi^{w}(x',x) \propto \exp\left(\frac{-\epsilon'\epsilon}{2}\right) \exp\left(\frac{\sigma W(x')}{2}\right)$$
$$\propto \exp\left(\frac{-\epsilon'\epsilon}{2}\right) \exp\left(\frac{-\sigma(\mu+C\epsilon)'P(\mu+C\epsilon)-\rho\sigma}{2}\right)$$
$$\propto \exp\left(\frac{-\epsilon'(I+\sigma C'PC)\epsilon}{2} - \epsilon'(I+\sigma C'PC)(I+\sigma C'PC)^{-1}\sigma C'P\mu\right)$$
$$\propto \exp\left[-.5\left(\epsilon+\sigma(I+\sigma C'PC)^{-1}C'P\mu\right)'(I+\sigma C'PC)\right.$$
$$\left(\epsilon+\sigma(I+\sigma C'PC)^{-1}C'P\mu\right)\right]$$

The last line portrays x' under π^w as having a Gaussian distribution with shocks that have mean vector $\tilde{\mu}$ and covariance matrix $\tilde{\Sigma}$ defined by

$$\tilde{\mu} = -\sigma (I + \sigma C' P C)^{-1} C' P \mu \qquad (17.2.9)$$

$$\tilde{\Sigma} = (I + \sigma C' P C)^{-1}.$$
 (17.2.10)

It follows that under the distorted conditional distribution, the mean and covariance matrix for x' are $A^*x + C\tilde{\mu}$ and $C\tilde{\Sigma}C'$. Notice that $-\sigma(I+\sigma C'PC)^{-1}C'P = \theta^{-1}(I-\theta^{-1}C'PC)^{-1}C'P\mu$, so that the formula for the distortion in the mean agrees with the distortion under the worst case model from chapters 2 and 6.⁸ In those chapters, we allowed the minimizing agent to distort only the mean,

 $^{^8\,}$ Mention small size of alteration in $\Sigma\,$ in HST example; mention continuous time no adjustment in limit.

not the variance, of the conditional distribution for next period's state. We have just shown, however, that when the minimizing agent is allowed to choose *any* distribution near the approximating transition density π , and when the density is Gaussian under the approximating model, the minimizing agent will select a Gaussian distorted distribution, but will choose to distort *both* the mean and the covariance matrix of the shocks. The formula for the mean distortion matches the one that prevails when the minimizing agent is allowed to alter only the mean vector.

17.3. Large deviation interpretation of \mathcal{R}

We have interpreted (17.2.7) in terms of a preference for robustness that is achieved by substituting the operator \mathcal{R} for the conditional expectations operator \mathcal{T} in a corresponding Bellman equation without a concern for robustness. In this section, we use ideas from the theory of large deviations to indicate how the operator $\mathcal{R}(W(x'))(x)$ contains information about the left tail of the distribution of W(x'). Recall from (17.2.2) that \mathcal{R} depends on α , and collapses to \mathcal{T} as $\alpha \nearrow +\infty$. We shall show that \mathcal{R} contains more information about the left tail of W(x') as α is decreased. We gather this interpretation from an exponential inequality that bounds the (conditional) tail probabilities of W. These tail probability bounds show how \mathcal{R} expresses a form of enhanced risk aversion that makes the decision-maker care about more than just the conditional mean of the continuation value.

The tail probability bound comes from the theory of large deviation approximations. It uses the inequality

$$\mathbf{1}_{\{W:W\leq -r\}}\leq \exp\left[\frac{-(W+r)}{\alpha}\right]$$

depicted in Figure 17.3.1. This inequality holds for any real number r and any $\alpha > 0$. Let z denote the state tomorrow. Then computing expectations conditioned on the current state vector x yields:

$$\Pr\{W(x') \le -r|x\} \le E\left(\exp\left[-\frac{W(x')}{\alpha}\right]|x\right)\exp\left(-\frac{r}{\alpha}\right),$$

or

$$\log[\Pr\{W(x') \le -r|x\}] \le -\frac{1}{\alpha}\mathcal{R}(W)(x) - \frac{r}{\alpha}.$$
(17.3.1)

The first term on the right side of this inequality is independent of r but depends on α . We can express (17.3.1) as

$$\Pr\{W(x') \le -r|x\} \le \exp\left\{-\alpha^{-1}\mathcal{R}(W)(x)\right\} \exp\left\{-\frac{r}{\alpha}\right\}.$$
 (17.3.2)

Inequality (17.3.2) bounds the tail probability on the left by an exponential in r. The right side declines with increases in r at rate $-\alpha^{-1}$; \mathcal{R} , a function of α , influences the constant in the bound. Decreasing α increases the exponential rate at which the bound sends the tail probabilities to zero, thereby expressing how a lower α heightens concern about tail events. This tells us how using \mathcal{R} to replace the mathematical expectation \mathcal{T} in a typical Bellman equation enhances risk aversion.



Figure 17.3.1: Ingredients of large deviation bounds: $\exp\left(\frac{-(W+r)}{\alpha}\right)$ and $\mathbf{1}_{\{W:W\leq -r\}}$ for r=1 and two values of α , 1 and 2.

Figure 17.3.2 shows how \mathcal{R} induces additional caution about continuation utilities W. In the figure, E(W) is the expected utility of a gamble between two continuation utility levels W_2, W_1 with $W_2 > W_1$. Where h(W) is a convex function, like $\exp(-W/\alpha)$ for $0 < \alpha < +\infty$, $h^{-1}E(h(W)) < E(W)$.



Figure 17.3.2: The function $h^{-1}E(h(W))$ for $h(W) = \exp(\frac{-W}{\alpha}), \ 0 < \alpha < +\infty.$

Chapter 18 Ramsey plans

18.1. Introduction

In chapter 16, a Stackelberg leader manipulates the worst-case model of the followers as part of its process of forming a Stackelberg plan. This chapter studies a Stackelberg leader in the person of a Ramsey planner who chooses both flat rate state contingent taxes and the system of date-state contingent prices to construct a competitive equilibrium that is best for a representative consumer. When the representative consumer fears model misspecification, the Ramsey planner chooses the consumer's worst-case belief in order to manipulate equilibrium state-date prices.

This chapter computes Ramsey equilibria for a version of Lucas and Stokey's (1983) model of taxation in an economy without physical capital in which the representative agent is concerned about misspecification of the model for government purchases and therefore appreciates robust decision rules. We modify Sargent and Velde's linear quadratic version of Lucas and Stokey's model.¹ In Lucas and Stokey's original model, the Ramsey planner manipulates both taxes and the price system in a way to maximize the utility of the representative agent subject to an implementability condition that requires that allocations are supported by competitive equilibria with flat rate taxes. Arrow-Debreu state-history prices can be represented as intertemporal rates of substitution times the probabilities that the representative agent assigns to transitions from one value the state to another. When the representative agent fears model misspecification, the transition probabilities that contribute to the Arrow-Debreu state-date prices correspond to the representative agent's worst-case model. In the Ramsey equilibrium with a representative consumer who fears misspecification, the Ramsey planner manipulates those worst-case probabilities as part of the process of choosing equilibrium prices. We begin by presenting the standard

– 405 –

¹ We thank Cristobal Huneeus and Yongseok Shin for excellent help with the computations. These François Velde's calculations in Sargent and Velde (1999XX).

model in which the representative agent is not concerned about misspecification.

18.1.1. Exogenous processes and information

Let x_t be an exogenous information vector. We shall use x_t to drive exogenous stochastic processes g_t, d_t, b_t, os_t , representing, respectively, government expenditures, an endowment, a preference shock, and a stream of promised coupon payments owed by the government at the beginning of time 0:

$$g_t = S_g x_t \tag{18.1.1a}$$

$$d_t = S_d x_t \tag{18.1.1b}$$

$$b_t = S_b x_t \tag{18.1.1c}$$

$$_{0}s_{t} = _{0}S_{s}x_{t}. \tag{18.1.1d}$$

Sargent and Velde (199XXX) made one of two alternative assumptions about the underlying stochastic process x_t .

Assumption 1: The $n \times 1$ vector process x_t with given initial condition x_0 is governed by

$$x_{t+1} = Ax_t + C\hat{\epsilon}_{t+1}.$$
 (18.1.2)

Here $\hat{\epsilon}_{t+1}$ is an i.i.d. random vector with mean 0 and identity covariance matrix, and A is a stable matrix.

Assumption 2: The process x_t is an n state Markov chain with time invariant transition probabilities arranged in the $n \times n$ matrix π with $\pi_{ij} = \operatorname{Prob}(x_{t+1} = \bar{x}_i | x_t = \bar{x}_i)$.

Later, when we find robust rules in section 18.2, we shall use one of the following assumptions about alternative specifications:

Assumption 1': The $n \times 1$ vector process x_t with given initial condition x_0 is governed by

$$x_{t+1} = Ax_t + C(\epsilon_{t+1} + w_{t+1}). \tag{18.1.3a}$$

Here ϵ_{t+1} is another i.i.d. random vector with mean 0 and identity covariance matrix, and w_{t+1} is a measurable function of $x_s, s \leq t$; w_{t+1} is a distortion to

the conditional mean of $\hat{\epsilon}_{t+1}$ in the approximating model (18.1.2). It satisfies

$$E_0 \sum_{t=0}^{\infty} \beta^{t+1} w_{t+1} \cdot w_{t+1} \le \eta_0.$$
(18.1.3b)

Assumption 2': The process x_t is one of a continuum of n state Markov chains indexed by matrices w with elements $w_{ij} > 0$. The transition probabilities are arranged in the $n \times n$ matrix π^w with $\pi^w_{ij} = (w_{ij}\pi_{ij}) / \sum_k w_{ik}\pi_{ik}$. Let $I(w)_i$ be called conditional entropy in state i, where

$$I(w)_i \equiv \sum_j \ln\left(\frac{w_{ij}}{\sum_k w_{ik}\pi_{ik}}\right) \left(\frac{w_{ij}\pi_{ij}}{\sum_k w_{ik}\pi_{ik}}\right).$$
 (18.1.4)

The constraint on misspecification is

$$E_0 \sum_{t=0}^{\infty} \beta^{t+1} I(w)_{t+1} \le \eta_0, \qquad (18.1.5)$$

where $I(w)_{t+1}$ is the value of conditional entropy when the system is in the *i*th state at time *t*.

In (18.1.4), the term $w_{ij}/\sum_k w_{ik}\pi_{ik}$ is the Radon-Nikodym derivative of the distorted transition density π^w with respect to the approximating π . Inequality (18.1.5) is a constraint on the entropy of Markov chain π^w relative to Markov chain π .²

We begin by staying with Sargent and Velde's assumptions 1 and 2, then adopt assumptions 1' and 2' when we study fear of misspecification.

 $^{^2}$ See Anderson, Hansen, and Sargent (2000) for a discussion of how to specify the set of alternative models when the approximating model is Markov.

18.1.2. Technology

There is a technology for converting one unit of labor ℓ_t into one unit of a single nonstorable consumption good. Feasible allocations satisfy:

$$c_t + g_t = d_t + \ell_t. \tag{18.1.6}$$

18.1.3. Representation of price system

From chapter 12, recall formula (12.3.2)

$$q_t(x_t|x_0) = \beta^t \frac{u'(c(x_t))}{e_1 \cdot u'(c(x_0))} f_t(x_t|x_0)$$

for the price vector at time t of a claim on list of consumption goods at time t when the state is x_t , where $f_t(x_t|x_0)$ is the representative consumer's transition t-step transition density for the state and u' is a vector marginal utilities of consumption. This equation simply rearranges the household's first-order for time-t, state x_t consumption. The present model has a single consumption good and the above formula becomes

$$q_t(x_t|x_0) = \beta^t \frac{b_t - c_t}{b_0 - c_0} f_t(x_t|x_0).$$
(18.1.7)

In the standard rational expectations model without concern about model misspecification, the same transition density $f_t(x_t|x_0)$ in this formula is accepted by both the government and the representative consumer. In the model in section 18.2 with concern about model misspecification, the government and the Ramsey planner will have a common f_t as their approximating model but different f_t 's corresponding to their worst-case models. When the representative consumer is concerned about model misspecification, we must replace (18.1.7) by

$$q_t(x_t|x_0) = \beta^t \frac{b_t - c_t}{b_0 - c_0} \hat{f}_t(x_t|x_0)$$
(18.1.8)

where $\hat{f}_t(x_t|x_0)$ is the worst-case *t*-step transition density of the representative consumer. In section 18.2, we shall provide an algorithm for computing the representative consumer's worst-case transition density $\hat{f}_t(x_t|x_0)$.

Introduction

As in chapter 12, we can follow Harrison and Kreps (1979XXXX) and define the scaled state-date price

$$p_t^0(x_t) = \frac{q_t(x_t|x_0)}{\beta^t f_t(x_t|x_0)},$$
(18.1.9)

which in the present context becomes

$$p_t^0(x_t) = \frac{b_t - c_t}{b_0 - c_0}.$$
(18.1.10)

18.1.4. Households

Markets are complete. At time 0, a representative consumer faces a scaled Arrow-Debreu price system $\{p_t^0\}$ and a flat rate tax on labor $\{\tau_t\}$ and chooses consumption and labor supply to maximize:

$$-.5E_0 \sum_{t=0}^{\infty} \beta^t [(c_t - b_t)^2 + \ell_t^2]$$
(18.1.11)

subject to the time 0 budget constraint

$$\sum_{t} \sum_{x_t} q_t(x_t | x_0) [d_t + (1 - \tau_t)\ell_t + {}_0s_t - c_t] = 0.$$
(18.1.12)

By using the scaled Arrow-Debreu prices, we can represent this equation in terms of the following conditional expectation:

$$E_0 \sum_{t=0}^{\infty} \beta^t p_t^0 [d_t + (1 - \tau_t)\ell_t + {}_0s_t - c_t] = 0.$$
 (18.1.13)

This states that the present value of consumption equals the present value of the endowment plus coupon payments on the initial government debt plus after-tax labor earnings. The scaled Arrow-Debreu price system is a stochastic process for which the conditional expectation is well defined.³

 $^{^3\,}$ See chapter 12 as well as Hansen (1987) and Hansen and Sargent (200XXX) for more about scaled Arrow-Debreu prices.

Ramsey plans

In section 18.2, we shall change (18.1.11) to reflect a concern about model misspecification by the representative consumer. As in chapter 12, when the representative agent fears model misspecification, it will be appropriate to use $\hat{f}_t(x_t|x_0)$ to define the scaled Arrow-Debreu prices in (18.1.9) and to replace E_0 in (18.1.13) with \hat{E} where \hat{E}_0 is the expectation with respect to the representative consumer's worst-case model.

18.1.5. Government

The government's time-0 budget constraint is

$$E_0 \sum_{t=0}^{\infty} \beta^t p_t^0 [(g_t + {}_0 s_t) - \tau_t \ell_t] = 0.$$
 (18.1.14)

Given the government expenditure process and the present value $E_0 \sum_{t=0}^{\infty} \beta^t p_t^0 s_t$, a budget-feasible tax process must satisfy (18.1.14).

18.1.6. Equilibrium

Definition: L_0^2 is the space of random variables y_t that are measurable with respect to x_t and such that $E_0 \sum_{t=0}^{\infty} \beta^t y_t^2 < +\infty$.

Definitions: A feasible allocation is a stochastic process $\{c_t, \ell_t\}$ that satisfies (18.1.6). A tax system is a scalar stochastic process $\{\tau_t\}$. A price system is a stochastic process $\{p_t^0\}$. The time t elements of each of these processes are assumed to be measurable with respect to x_t and to belong to L_0^2 .

Definition: An *equilibrium* is a feasible allocation, a price system, and a tax system that have the following properties:

i. Given the tax and price systems, the allocation solves the household's problem.

ii. Given the price system, the allocation and the tax system satisfy the government's budget constraint.

18.1.7. Properties

The first-order conditions for the household's problem imply that the equilibrium price system satisfies $p_t^0 = \xi(b_t - c_t)$, where ξ is a numeraire that we set at $(b_0 - c_0)^{-1}$. As in chapter 12, the preference specification permits the scaled Arrow-Debreu price p_t^0 to be expressed in terms of ratios of linear functions of the state:

$$p_t^0 = M_p x_t / M_p x_0,$$

where M_p is a matrix defined so that $M_p x_t = b_t - c_t$. The preference specification will make it possible to express government time t revenues as the ratio of a quadratic function of the state at t to a linear function of the state at 0. The forms of these prices and taxes and of the other objects in (18.1.11) reduce the technical problem to evaluating geometric sums of a quadratic form in the state, which means solving Sylvester equations. For assumptions 1 and 2, Appendix A of Sargent and Velde (1999) shows how to compute such sums by using standard formulas for expectations of geometric sums of a quadratic form.

18.1.8. Ramsey problem

There are many equilibria, indexed by tax systems. The Ramsey problem is to choose the tax system that delivers the equilibrium preferred by the representative household. The Ramsey problem assumes that at time 0 the government commits itself to the tax system, once and for all.

Definition: The Ramsey problem is to choose an equilibrium that maximizes the household's welfare (18.1.11). The allocation that solves this problem is called the Ramsey allocation, and the associated tax system is called the Ramsey plan.

18.1.9. Solution strategy

Following a long line of researchers starting with Frank Ramsey (1929), we shall solve this problem using a 'first-order' approach that involves the following steps. The steps incorporate the properties required by the definition of equilibrium.

1. Obtain the first-order conditions for the household's problem and use them to express the tax system and the price system in terms of the allocation.

2. Substitute the expressions for the tax system and the price system obtained in step 1 into the government's budget constraint to obtain a single intertemporal restriction on allocations.

3. Use Lagrangian methods to find the feasible allocation that maximizes the utility of the representative household subject to the restriction derived in step 2. The maximizer is the Ramsey allocation.

4. Use the expressions from step 1 to find the associated Ramsey equilibrium price and tax systems by evaluating them at the Ramsey allocation.

18.1.10. Computation with no concern about robustness

We now execute these four steps for the version of the model without a preference for robustness. The problem is set so that the mathematics of linear systems can support a solution.

Step 1. The household's first order conditions imply

$$p_t^0 = (b_t - c_t)/(b_0 - c_0)$$
(18.1.15)

$$\tau_t = 1 - \frac{\ell_t}{b_t - c_t}.$$
(18.1.16)

Step 2. Using (18.1.15) and (18.1.16), express (18.1.14) as

$$E_0 \sum_{t=0}^{\infty} \beta^t [(b_t - c_t)(g_t + {}_0s_t) - (b_t - c_t)\ell_t + \ell_t^2] = 0.$$
(18.1.17)

Equation (18.1.17) is often called the implementability constraint on the allocation. Step 3. Consider the maximization problem associated with the Lagrangian:

$$J = E_0 \sum_{t=0}^{\infty} \beta^t \left\{ -.5[(c_t - b_t)^2 + \ell_t^2] + \lambda_0[(b_t - c_t)\ell_t - \ell_t^2 - (b_t - c_t)(g_t + {}_0s_t)] + \mu_{0t}[d_t + \ell_t - c_t - g_t] \right\}$$

where λ_0 is the multiplier associated with the government's budget constraint, and μ_{0t} is the multiplier associated with the time t feasibility condition. Obtain the first-order conditions:

$$c_t: -(c_t - b_t) + \lambda_0 [-\ell_t + (g_t + {}_0s_t)] = \mu_{0t}$$
(18.1.18a)

$$\ell_t : \ell_t - \lambda_0 [(b_t - c_t) - 2\ell_t] = \mu_{0t}$$
(18.1.18b)

$$\mu_{0t}: \quad d_t + \ell_t = c_t + g_t \tag{18.1.18c}$$

We want to solve equations (18.1.18a), (18.1.18b), (18.1.18c) and the government's budget constraint (18.1.14) for an *allocation*.

18.1.11. Key idea

Our solution strategy is to begin by taking λ_0 as given and to solve (18.1.18) for an allocation contingent on λ_0 . Then we shall use (18.1.14) to solve for λ_0 .

18.1.12. Execution

Using the feasibility constraint $c_t = d_t + \ell_t - g_t$, we can express (18.1.18*a*), (18.1.18*b*) as

$$\ell_t - \lambda_0 [(b_t - d_t - \ell_t + g_t) - 2\ell_t] = -(d_t + \ell_t - g_t - b_t) + \lambda_0 [-\ell_t + (g_t + 0s_t)]$$

or

$$\ell_t = \frac{1}{2}(b_t - d_t + g_t) - \frac{\lambda_0}{2 + 4\lambda_0}(b_t - d_t - {}_0s_t).$$

We also derive

$$c_t = \frac{1}{2}(b_t + d_t - g_t) - \frac{\lambda_0}{2 + 4\lambda_0}(b_t - d_t - {}_0s_t).$$

Define

$$\tilde{c}_t = (b_t + d_t - g_t)/2 \tag{18.1.19a}$$

$$\tilde{\ell}_t = (b_t - d_t + g_t)/2 \tag{18.1.19b}$$

$$m_t = (b_t - d_t - {}_0s_t)/2 \tag{18.1.19c.}$$

We have:

$$\ell_t = \ell_t - \mu m_t \tag{18.1.20a}$$

$$c_t = \tilde{c}_t - \mu m_t \tag{18.1.20b}$$

where, for convenience, we define

$$\mu = \frac{\lambda_0}{1 + 2\lambda_0}.$$
 (18.1.21)

Using (18.1.20), the general term of (18.1.17) can be written as:

$$\begin{aligned} (b_t - \tilde{c}_t)(g_t + {}_0s_t) - (b_t - \tilde{c}_t)\tilde{\ell}_t + \tilde{\ell}_t^2 \\ &- \mu m_t [-(g_t + {}_0s_t) + \tilde{\ell}_t - (b_t - \tilde{c}_t) + 2\tilde{\ell}_t] + \mu^2 m_t^2 \\ &= (b_t - \tilde{c}_t)(g_t + {}_0s_t) - 2m_t^2 \mu + 2m_t^2 \mu^2, \end{aligned}$$

where we used $2\tilde{\ell}_t = b_t - d_t + g_t$ and $\tilde{\ell}_t = b_t - \tilde{c}_t$ to reduce the bracketed factor in the second line.

This allows us to write (18.1.17) as:

$$a_0(x_0)(\mu^2 - \mu) + b_0(x_0) = 0$$
(18.1.22)

where

$$a_{0}(x_{0}) = E_{0} \sum_{t=0}^{\infty} \beta^{t} \frac{1}{2} (b_{t} - d_{t} - {}_{0}s_{t})^{2}$$

= $E_{0} \sum_{t=0}^{\infty} \beta^{t} x_{t}' \frac{1}{2} [S_{b} - S_{d} - {}_{0}S_{s}]' [S_{b} - S_{d} - {}_{0}S_{s}] x_{t}$ (18.1.23)

and

$$b_0(x_0) = E_0 \sum_{t=0}^{\infty} \beta^t [(b_t - \tilde{c}_t)(g_t + {}_0s_t) - (b_t - \tilde{c}_t)\tilde{\ell}_t + \tilde{\ell}_t^2] \quad (18.1.24)$$
$$= E_0 \sum_{t=0}^{\infty} \beta^t \frac{1}{2} (b_t - d_t + g_t)(g_t + {}_0s_t)$$
$$= E_0 \sum_{t=0}^{\infty} \beta^t \frac{1}{2} x_t' [S_b - S_d + S_g]' [S_g + {}_0S_s] x_t, \quad (18.1.25)$$

414

where we have used the fact that $b_t - \tilde{c}_t = \tilde{\ell}_t$. The 0 subscripts on the quadratic forms a_0 and b_0 denote their dependence on ${}_0S_s$. The coefficients in the polynomial expression (18.1.22) are functions of x_0 alone because, given the law of motion for the exogenous state x_t , the infinite sums can be computed using the algorithms described in Appendix A of Sargent and Velde (1999).

Notice that $b_0(x_0)$, when expressed by (18.1.24), is simply the infinite sum on the left side of (18.1.17) evaluated for the specific allocation $\{\tilde{c}_t, \tilde{\ell}_t\}$ defined in (18.1.19), which solves the problem:

$$\max_{c,\ell} -.5[(c-b_t)^2 + \ell^2]$$

subject to $c + g_t = \ell + d_t$; $\{\tilde{c}_t, \tilde{\ell}_t\}$ is the allocation that would be chosen by a benevolent dictator able to choose among *all* feasible allocations, not just competitive equilibrium allocations. This is also the Ramsey allocation when the government can resort to lump-sum taxation. The term $b_0(x_0)$ is the presentvalue of the government's stream of spending commitments $\{g_t+_0s_t\}$, evaluated at the prices that correspond to the $\{\tilde{c}_t, \tilde{\ell}_t\}$ allocation. If that present value is 0, distortionary taxation is not necessary, and $\mu = 0$ (which implies that $\lambda_0 = 0$) solves (18.1.22): the government's budget constraint is not binding. One configuration for which $b_0(x_0) = 0$ occurs when $g_t = -_0s_t$ for all t, but there are many others. Because markets are complete, the timing of the government's claims on the household does not matter. If the government were able to acquire such claims on the private sector in a non-distortionary way, it would be able to implement a first-best allocation.

When the net present value of the government's commitments is positive, we must solve (18.1.22) for a μ in (0,1/2) that correspond to a $\lambda_0 > 0$. The polynomial $a_0(x_0)\mu(1-\mu)$ is bounded above by $a_0(x_0)/4$, which means that government commitments that are "too large" cannot be supported by a Ramsey plan. If $b_0(x_0) < a_0(x_0)/4$, there exists a unique solution μ in (0,1/2) and a unique $\lambda_0 > 0$. The Ramsey allocation can then be computed as:

$$c_t = \tilde{c}_t - \mu m_t$$

= $\frac{1}{2}([S_b + S_d - S_g] - \mu[S_b - S_d - {}_0S_s])x_t$ (18.1.26a)

$$\ell_t = \ell_t - \mu m_t$$

= $\frac{1}{2} ([S_b - S_d + S_g] - \mu [S_b - S_d - {}_0S_s]) x_t$ (18.1.26b)

and the Ramsey plan as:

$$\tau_{t} = 1 - \frac{\ell_{t}}{b_{t} - c_{t}}$$

$$= 1 - \frac{\tilde{\ell}_{t} - \mu m_{t}}{b_{t} - \tilde{c}_{t} + \mu m_{t}}$$

$$= \frac{2\mu m_{t}}{\tilde{\ell}_{t} + \mu m_{t}}$$

$$= \frac{2\mu [S_{b} - S_{d} - {}_{0}S_{s}]x_{t}}{([S_{b} - S_{d} + S_{g}] + \mu [S_{b} - S_{d} - {}_{0}S_{s}])x_{t}}.$$
(18.1.27)

Expression (18.1.27) shows that when the endowment and the preference shocks are constant, the stochastic properties of the tax rate mirror those for government expenditures. Tax rates vary inversely with government expenditures (notice that $S_g x_t$ appears in the denominator).

18.2. Modifications for robustness

We now compute a Ramsey plan when the representative consumer is concerned about model misspecifications of a type described either by Assumption 1' or by Assumption 2'. Now the $\hat{f}_t(x_t|x_0)$ that occurs in the equilibrium state-date prices $q_t(x_t|x_0) = \beta^t \frac{b_t - c_t}{b_0 - c_0} \hat{f}_t(x_t|x_0)$ defined in (18.1.8) becomes the representative consumer's worst-case transition density. To compute the Ramsey plan, we must use these modified $q_t(x_t|x_0)$'s to evaluate the government's budget constraint. To deduce \hat{f}_t , we use a multiplier problem with parameter θ for the representative consumer at a candidate Ramsey allocation. We let θ be a multiplier on the constraint on the specification error (18.1.3b) for Assumption 1' or (18.1.5) for Assumption 2'. We take $\theta \in (\underline{\theta}, +\infty]$ as a parameter. The value of $\underline{\theta}$ is context specific because it depends on the government expenditure process.⁴

⁴ See chapter 7 for a discussion of the lower limit on θ .

18.2.1. Main idea

Recall our basic solution strategy of (1) taking λ_0 as given, (2) solving (18.1.18) for an allocation contingent on λ_0 , then (3) using (18.1.14) to solve for λ_0 . We solved the Ramsey problem by searching for a λ_0 that solves (18.1.14). For computing a Ramsey plan when the representative consumer fears model misspecification, steps (1) and (2) remain the same, and (3) is modified by the extra step of finding a distorted expectations operator with which to evaluate the government's budget constraint (18.1.14):

$$\hat{E}_0 \sum_{t=0}^{\infty} \beta^t p_t^0 [(g_t + {}_0 s_t) - \tau_t \ell_t] = 0$$
(18.2.1)

where \hat{E} is evaluated with respect to \hat{f}_t . We describe how to compute this modified expectations operator under Assumptions 1' and 2'.

18.2.2. Assumption 1'

We first adopt Assumption 1' and form the multiplier problem for the representative consumer at a candidate Ramsey allocation associated with a fixed λ_0 . For a given λ_0 , the candidate allocation is held fixed, so that this becomes a pure minimization problem. It is important at this point to observe that given \hat{f}_t , the first-order conditions for the Ramsey planner retain the same form (18.1.19), (18.1.20) that they took in the original model without a concern about robustness on the part of the representative consumer. These equations lead directly to the formulas (18.1.26) for the Ramsey allocation in terms of the multiplier μ and the state x_t . The adjustment for robustness occurs entirely through the transformed multiplier μ due to replacing f_t from the approximating model with \hat{f}_t from the consumer's worst-case model in evaluating the government's budget constraint (18.2.1).

Thus, we proceed as follows. Given a candidate λ_0 , we use (18.1.26), (18.1.1), and (18.1.11) to form the matrix H in the representation $x'_t H' H x_t = -.5[(b_t - c_t)^2 + \ell_t^2]$ of the representative consumer's one-period utility function at the candidate Ramsey allocation associated with a given candidate μ . Then to deduce the worst-case model to be used to compute the distorted transition density needed to form the Arrow-Debreu prices and the budget constraint

(18.2.1), we form the multiplier problem:

$$\min_{w} E_0 \sum_{t=0}^{\infty} \beta^t \left\{ -x'_t H' H x_t + \beta \theta w'_{t+1} w_{t+1} \right\}$$
(18.2.2)

where the expectation and the minimization are both subject to

$$x_{t+1} = Ax_t + C(\epsilon_{t+1} + w_{t+1}).$$
(18.2.3)

This is a discounted optimal linear regulator problem with optimal feedback rule

$$w_{t+1} = Kx_t. (18.2.4)$$

Substituting (18.2.4) into (18.2.3) gives the distorted law of motion

$$x_{t+1} = \hat{A}x_t + C\epsilon_{t+1} \tag{18.2.5}$$

where

$$A = A + CK. \tag{18.2.6}$$

An alternative formula for K can be found as follows. First, compute a matrix V by iterating to convergence on

$$\mathcal{S}(V_{i+1}) = -H'H + \beta A'\mathcal{D}(V_i)A$$

where \mathcal{D} is the operator

$$\mathcal{D}(V) = V + \sigma V C (I - \sigma C' V C)^{-1} C' V$$
(18.2.7)

and $\sigma \equiv -\theta^{-1}$. Then compute K from

$$K = \sigma (I - \sigma C' V C)^{-1} C' V. \qquad (18.2.8)$$

The fact that H depends on μ , or λ_0 , through (18.1.19), (18.1.20) necessitates that now we use an iterative method to compute λ_0 , whereas Sargent and Velde could calculate it directly.

Here is a four step process to compute the Ramsey plan when the representative consumer's concern about robustness is associated with given θ :

Step 1. Guess a value of λ_0 . Find the associated μ from (18.1.21).

418

Step 2. Compute c_t, ℓ_t from (18.1.19), (18.1.20). Find the associated H for (18.2.2). Find the associated K and \hat{A} .

Step 3. Using model (18.2.5), (18.2.6) to evaluate the expectation operator \hat{E} , evaluate $a_0(x_0)$ and $b_0(x_0)$ using the formulas

$$a_0(x_0) = \hat{E}_0 \sum_{t=0}^{\infty} \beta^t x_t' \frac{1}{2} [S_b - S_d - {}_0S_s]' [S_b - S_d - {}_0S_s] x_t$$
(18.2.9)

and

$$b_0(x_0) = \hat{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{1}{2} x_t' [S_b - S_d + S_g]' [S_g + {}_0S_s] x_t.$$
(18.2.10)

Step 4. Check whether (18.1.22) is satisfied. Iterate on steps 1 through 3 to find a λ_0 that is a zero of (18.1.22).

Step 5. Having found such a λ_0 , use (18.1.26) and (18.1.27) to compute the Ramsey allocation and Ramsey plan.

18.2.3. Assumption 2'

Here is the corresponding step-by-step plan for computing the Ramsey plan when the representative consumer is concerned about robustness under the Markov assumption 2'. First, again fix a λ_0 and use (18.1.26), (18.1.1), and (18.1.11) to define H in the following representation of the one-period return function:

$$u(x_t) = -x'_t H' H x_t. (18.2.11)$$

This is the same H as found above. Then follow these steps.

Step 1. Compute the value function V (a vector) by iterating to convergence on

$$V_i = u_i - \beta \theta \ln \left\{ \sum_j \exp\left(\frac{-V_j}{\theta}\right) \pi_{ij} \right\}.$$
 (18.2.12)

Step 2. Form

$$w_j^* = \exp\left(\frac{-V_j}{\theta}\right). \tag{18.2.13}$$

Step 3. Form the distorted transition density

$$\pi_{ij}^{w^*} = \frac{w_j^* \pi_{ij}}{\sum_k w_k^* \pi_{ik}}.$$
(18.2.14)

Step 4. Evaluate $a_0(x_0), b_0(x_0)$, using π^{w^*} and formulas for computing geometric sums of a quadratic form for a Markov chain.⁵

Step 5. Check whether (18.1.22) is satisfied. Iterate on steps 1 through 4 to find a λ_0 that is a zero of (18.1.22).

Step 6. Having found such a λ_0 , use (18.1.26) and (18.1.27) to compute the Ramsey allocation and Ramsey plan.

18.2.4. Ramsey planner's worst-case model

So far, we have concentrated entirely on the representative consumer's attitudes about model misspecification and the representative consumer's worstcase model. In this section we explain why the essentially static nature of the Ramsey problem mean that the planner's own worst-case model affects neither the Ramsey plan nor the Ramsey allocation.

The indirect preferences of the Ramsey planner differ from those of the representative consumer because the Ramsey planner must respect the implementability condition that restricts him to competitive equilibrium allocations.⁶ Therefore if we were to attribute a concern about model misspecification to the Ramsey planner, the worst-case model of the Ramsey planner would typically differ from the worst-case model of the representative consumer.⁷ However, given λ_0 , the first-order conditions for the maximization part of the Ramsey planner's problem are entirely static. That fact implies that the Ramsey planner's worst-case model has no affect on the Ramsey plan or on the Ramsey allocation. The Ramsey plan and the Ramsey allocation can be computed without

⁵ See Sargent and Velde (1999).

⁶ See Jones, Manuelli, and Rossi (19XX) for an interpretation of the Ramsey planners problem as being an ordinary planning problem with a weird one-period utility function that is augmented by the time-t component of the implementability restriction.

⁷ Recall the characterization of time-inconsistency in chapter 16, page 376, in terms of conflict or misalignment of preferences.
computing the planner's worst-case model.⁸ The worst-case model of the representative consumer is what gets embedded in equilibrium prices. The Ramsey planner affects equilibrium prices partly by manipulating the worst-case model of the representative consumer.

18.3. Computations

We activate a concern about robustness on the part of the representative consumer by setting values of $\theta < +\infty$ in (18.2.2) or (18.2.12). We want to study how allocations, taxes, and prices change as we accentuate a preference for robustness (i.e., lower θ) for both the 1' and 2' cases. In addition to these objects, we also want to calculate government debt and interest rates. Along the Ramsey allocation, government debt B_t equals

$$B_t = \frac{\hat{E}_t \sum_{j=0}^{\infty} \beta^j [(b_{t+j} - c_{t+j})\ell_{t+j} - \ell_{t+j}^2 - (b_{t+j} - c_{t+j})g_{t+j}]}{(b_t - c_t)}, \quad (18.3.1)$$

which can evidently be expressed as a function of the time t state x_t , in particular, a quadratic form in x_t plus a constant divided by a linear form in x_t . The quantity B_t can be regarded as the time t value of government state contingent debt (Arrow securities) issued at t-1. We also compute the one period gross interest rate from

$$R_t^{-1} = \frac{\dot{E}_t \beta p_{t+1}^t}{p_t^t},$$

where $p_{t+1}^t = \frac{M_p x_{t+1}}{M_p x_t}$.

⁸ Computing the Ramsey planner's worst-case model amounts to solving a 'pure prediction problem' using the weird one-period preferences described in the previous footnote.

18.4. Simulations at approximating model

Using the approximating model to generate the government expenditure process, we simulate Ramsey plans with a representative consumer who fears model misspecification and compare them with Ramsey plans for a representative consummer who trusts the approximating model. Under assumption 1', for an approximating model A, C and a given θ , a Ramsev plan with robust consumers is affiliated with a distorted law of motion A, C that is the worst-case model of the representative agent and that gets embedded in equilibrium state-date prices. At the approximating model, we are interested in comparing the Ramsey plan with a robust representative consumer with the plan for representative consumer who completely trusts the model. We simulate systems with different θ 's at the approximating model (i.e., A not \hat{A}). These simulations thus indicate outcomes when the approximating model for the government expenditure process is indeed correct. Similarly, for Assumption 2', we simulate the robust Ramsey plan using the approximating Markov chain π to generate government expenditures, while using the representative consumer's distorted chain π^{w*} to evaluate asset prices and budget constraints.

18.5. Computed Ramsey plans and allocations

18.5.1. Markov case

We specify a three state Markov chain with transition matrix

$$\pi = \begin{bmatrix} .95 & .05 & 0\\ 0 & .8 & .2\\ 0 & 0 & 1 \end{bmatrix},$$

and initial distribution $\pi_0 = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}'$, so that the system starts in state 1. Government expenditures in states 1, 2, 3 are $\begin{bmatrix} .5 & .5 & .25 \end{bmatrix}$. State 1 is designed to represent war, state 2 armistice, and state 3 peace. Peace is an absorbing state. Achieving peace requires spending at least one period in the state of armistice. We specify that $\beta = .97$ and b = 2.2.

For two values of θ , namely, 4 and 1, Fig. 18.5.1 and Fig. 18.5.2 show the Ramsey plans and allocations with (solid lines) and without (dotted lines) a

concern for robustness on the part of the representative consumer. Recall that lowering θ raises the preference for robustness. For both values of θ , we have computed the associated distorted Markov transition matrices. With $\theta = 4$ we have

$$\pi^{w^*} = \begin{bmatrix} .9836 & .0164 & 0 \\ 0 & .8581 & .1419 \\ 0 & 0 & 1 \end{bmatrix}.$$

With $\theta = 1$ we have

$$\pi^{w^*} = \begin{bmatrix} .99975 & .00025 & 0 \\ 0 & .99393 & .00607 \\ 0 & 0 & 1 \end{bmatrix}.$$

Notice how raising the consumer's concern about model misspecification (i.e., lowering θ) shifts probability toward longer wars and a longer armistice.

Relative to the standard Ramsey plan where the representative consumer trusts the model, the Ramsey plan where the consumer is concerned about robustness has taxes higher and consumption and labor supply both lower. Once peace occurs, taxes and government debt are both higher in the robust Ramsey plan. After peace arrives, the representative consumer's worst-case beliefs coincide with the approximating beliefs. Thus, the higher taxes and government debt during peace time in the robust plan reflect the government's having to honor its past promises, entered into during earlier periods of war or armistice when it was facing state-date prices that embedded the representative consumer's worstcase beliefs via formula (18.1.8). During war and peace, interest rates coincide for the Ramsey plans with and without a concern for robustness by the representative consumer, however, they are lower during the armistice when the representative consumer wants robustness. This reflects the consumer's more pessimistic expectation about the one-period rate of growth of consumption under the distorted Markov chain π^{w^*} used to price assets.⁹

When we push θ down to 1, the government sets taxes so high that the value of its debt is negative at first. This is the government's response to the twisted state-date prices it faces due to the representative consumer's pessimism. At those prices, the government actually acquires claims on the public that it uses to fund some of its expenditures during armistice.

 $^{^9\,}$ Notice that during war the expected rate of change of consumption under π and π^{w*} coincide.



Figure 18.5.1: Markov case, $\theta = 4$. Dotted line is the Ramsey plan without a preference for robustness, the solid line is the Ramsey plan with a preference for robustness.

18.5.2. Stochastic difference equation

Fig. 18.5.3 compares Ramsey plans with and without a concern about robustness on the part of the representative consumer when government expenditures follow the first order autoregression:

$$x_{t+1} = .0175 + .95x_t + C\epsilon_{t+1},$$

where C = .013660. When $\theta = .1$, the associated worst-case distorted law of motion for government expenditures is

$$x_{t+1} = .0476 + .9580x_t + C\epsilon_{t+1}$$



Figure 18.5.2: Markov case, $\theta = 1$. Dotted line is the Ramsey plan without robustness, the solid line is the Ramsey plan with a representative consumer who wants robustness.

where again $C = .013660.^{10}$ The mean of x_t under A is .35 while under the distorted law it is 1.1324. Thus, pessimism translates into more persistence and a higher mean for government expenditures. We set the other parameters at $\beta = .97$ and b = 2.135.

Fig. 18.5.3 shows taxes and government debt to be higher, and labor and consumption lower when the representative consumer wants robustness. Interest rates are uniformly lower when the consumer wants robustness, reflecting the consumer's pessimism about the rate of growth of consumption embedded in the distorted law of motion \hat{A} .

 $^{^{10}\,}$ We have set θ to a low value to accentuate the effects of robustness so that they show up well on the graphs.



Figure 18.5.3: Stochastic difference equation case, $\theta = .1$. Dotted line is the Ramsey plan without robustness, the solid line is the Ramsey plan with a representative consumer who wants robustness.

References

- Anderson, B. D. O. (1978). 'Second-order Convergent Algorithms for the Steady-state Riccati Equation'. International Journal of Control, Vol. 28, pp. 295–306.
- Anderson, B. D. O. and J. B. Moore (1979). *Optimal Filtering*. Englewood Cliffs, NJ: Prentice Hall.
- Anderson, E. (1995). 'Computing Equilibria in Linear-Quadratic Dynamic Games and Models with Distortions'. Mimeo. University of Chicago.
- Anderson, E. (1998). 'Uncertainty and the Dynamics of Pareto Optimal Allocations'. Mimeo. University of Chicago Dissertation.
- Anderson, E., L. P. Hansen, E. R. McGrattan, and T. J. Sargent (1996). 'Mechanics of Forming and Estimating Dynamic Linear Economies'. In H. Amman, D. A. Kendrick and J. Rust (eds.), Handbook of Computational Economics, Vol.1. Amsterdam: North Holland.
- Anderson, E., L. P. Hansen and T. J. Sargent (1999). 'Robustness, Detection and the Price of Risk'. Mimeo. March, ???? .
- Anderson, G. and G. Moore (1985). 'A Linear Algebraic Procedure for Solving Linear Perfect Foresight Models'. *Economics Letters*, Vol. 17, pp. 247–252.
- Ansley, C. F. and R. Kohn (1985). 'Estimation, Filtering, and Smoothing in State-Space Models with Incompletely Specified Initial Conditions'. Annals of Statistics, Vol. 13, pp. 1286–1316.
- Apostol, T. M. (1975). Mathematical Analysis. 2nd ed. Reading, MA: Addison-Wesley.
- Backus, David and John Driffill (1986). 'The Consistency of Optimal Policy in Stochastic Rational Expectations Models'. Mimeo. CEPR discussion paper No. 124.
- Başar, T. and P. Bernhard (1995). H^{∞} -Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach. Boston-Basel-Berlin: Birkhäuser.
- Bai, Z. and J. W. Demmel (1993). 'On Swapping Diagonal Blocks in Real Schur Form'. Linear Algebra and its Applications, Vol. 186, pp. 73–95.
- Bailey, Martin (1971). National Income and the Price Level, Second Edition. New York, McGraw Hill, pp. 175-186.
- Ball, L. (1999). 'Policy Rules for Open Economies'. In J. B. Taylor (ed.), Monetary Policy Rules University of Chicago Press, pp. 127–144.
- Barsky, R. T., T. Juster, M. Kimball and M. Shapiro (1997). 'Preference Parameters and Behavior Heterogeneity: An Experimental Approach in the Health and Retirement Survey'. Quarterly Journal of Economics, Vol. 12, pp. 537–580.
- Bartels, R. H. and G. W. Stewart (1972). 'Algorithm 432 Solution of the Matrix Equation AX + XB = C'. Communications of the ACM, Vol. 15, pp. 820–826.
- Basar, T. and P. Bernhard (1995). H[∞]-Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach. Boston, MA: Birkhauser.
- Becker, G. S. and K. M. Murphy (1988). 'A Theory of Rational Addiction'. Journal of Political Economy, Vol. 96, pp. 675–700.
- Bierman, G. J. (1984). Computational Aspects of the Matrix Sign Function Solution to the ARE. Proceedings 23rd IEEE Conference on Decision Control, pp. 514–519

- 427 -

- Blanchard, O. J. and C. M. Kahn (1980). 'The Solution of Linear Difference Models under Rational Expectations'. *Econometrica*, Vol. 48(5), pp. 1305–1311.
- Blinder, A. S. (1998). Central Banking in Theory and Practice. Cambridge, MA: MIT Press.
- Brainard, W. (1967). 'Uncertainty and the Effectiveness of Policy'. American Economic Review, Vol. 57, pp. 411–425.
- Bray, M. B. and D. M. Kreps (1987). 'Rational Learning and Rational Expectations'. In G. R. Feiwel (ed.), Arrow and the Ascent of Modern Economic Theory New York University Press, pp. 597–625.
- Brock, W. A. and L. Mirman (1972). 'Optimal Economic Growth and Uncertainty: The Discounted Case'. Journal of Economic Theory, Vol. 4(3), pp. 479–513.
- Brunner, K. and A. Meltzer (1969). 'The Nature of the Policy Problem'. In K. Bruner and A. Meltzer (eds.), *Targets and Indicators of Monetary Policy*. Chandler Publishing Company, San Francisco.
- Byers, R. (1987). 'Solving the Algebraic Riccati Equation with the Matrix Sign Function'. Linear Algebra and its Applications, Vol. 85, pp. 267–279.
- Cagetti, M., L. P. Hansen, T. J. Sargent and N. Williams (2000). 'Robustness and Filtering in a Stochastic Growth Model'. Mimeo. Minnesota Macro Conference.
- Caines, P. E. (1988). Linear Stochastic Systems. New York: John Wiley & Sons, Inc..
- Caines, P. E. and D. Q. Mayne (1970). 'On the Discrete Time Matrix Riccati Equation of Optimal Control'. International Journal of Control, Vol. 12(5), pp. 785–794.
- Caines, P. E. and D. Q. Mayne (1971). 'Correspondence: "On the Discrete Time Matrix Riccati Equation of Optimal Control-a Correction". International Journal of Control, Vol. 14(1), pp. 205–207.
- Camerer, C. (1995). 'Individual Decision Making'. The Handbook of Experimental Economics, edited by J. H. Kagel and A. Roth, pp. 588–673.
- Chan, S. W., G. C. Goodwin and K. S. Sin (1984). 'Convergence Properties of the Riccati Difference Equation in Optimal Filtering of Nonstabilizable Systems'. *IEEE Transactions on Automatic Control*, Vol. AC-29(2), pp. 110–118.
- Chen, Z. and L. Epstein (1998). 'Ambiguity, Risk, and Asset Returns in Continuous Time'. Mimeo. August, ????? .
- Chernoff, H. (1952). 'A Measure of Asymptotic Efficiency for Tests of a Hypothesis Based on Sums of Observations'. Annals of Mathematical Statistics, Vol. 23, pp. 493–507.
- Christiano, L. J. and T. J. Fitzgerald (1998). 'The Business Cycle: Its Still a Puzzle'. Federal Reserve Bank of Chicago Economic Perspectives, Vol. 4, pp. 56–83.
- Christiano, L. J. and C. J. Gust (1999). 'Comment'. In J. B. Taylor (ed.), Monetary Policy Rules Chicago: University of Chicago Press, pp. 299–316.
- Clarida, R., J. Gali and M. Gertler (1999). 'The Science of Monetary Policy: A New Keynesian Perspective'. Journal of Economic Literature, Vol. 37(4), pp. 1661–1707.
- Clarida, R., J. Gali and M. Gertler (2000). 'Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory'. Quarterly Journal of Economics, Vol. 115(1), pp. 147–180.
- Cochrane, J. H. and L. P. Hansen (1992). 'Asset Pricing Explorations for Macroeconomics'. NBER Macroeconomics Annual, pp. 115–169.

- Currie, D. and P. Levine (1987). 'The design of feedback rules in linear stochastic rational expectations models'. Journal of Economic Dynamics and Control, Vol. 11(1), pp. 1–28.
- Dantzig, George B. (1998). Linear Programming and Extensions. Eleventh printing, Princeton University Press.
- De Jong, D., B. Ingram and C. Whiteman (1996). 'A Bayesian Approach to Calibration'. Journal of Business and Economic Statistics, Vol. 14, January, pp. 1–10.
- Dempster, A. (1967). 'Upper and Lower Probabilities Induced by a Multivalued Map'. Annals of Mathematical Statistics, Vol. 38, pp. 325–339.
- Denman, E. D. and A. N. Beavers (1976). 'The Matrix Sign Function and Computations in Systems'. Applications of Mathematical Computations, Vol. 2, pp. 63–94.
- Doan, T., R. Litterman and C. Sims (1984). 'Forecasting and Conditional Projection Using Realistic Prior Distributions'. *Econometric Reviews*, Vol. 3(1), pp. 1–100.
- Dow, J. and S. Werlang (1994). 'Learning under Knightian Uncertainty: The Law of Large Numbers for Non-additive Probabilities'. Mimeo. ????? .
- Dupuis, P. and R. S. Ellis (1997). A Weak Convergence Approach to the Theory of Large Deviations. New York: John Wiley & Sons.
- Economic and Statistics Administration, and Bureau of the Census (1993). The American Almanac 1993-1994. Austin: The Reference Press.
- Ellsberg, D. (1961). 'Risk, Ambiguity and the Savage Axioms'. Quarterly Journal of Economics, Vol. 75, pp. 643–669.
- Epstein, L. G. and A. Melino (1995). 'A Revealed Preference Analysis of Asset Pricing Under Recursive Utility'. Review of Economic Studies, Vol. 62, pp. 597–618.
- Epstein, L. G. and T. Wang (1994). 'Intertemporal Asset Pricing Under Knightian Uncertainty'. Econometrica, Vol. 62(3), pp. 283–322.
- Epstein, L. G. and S. E. Zin (1989). 'Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework'. *Econometrica*, Vol. 57(4), pp. 937–969.
- Epstein, L. G. and S. E. Zin (1991). 'Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis'. Journal of Political Economy, Vol. 99(2), pp. 263–286.
- Fellner, W. (1961). 'Distortion of Subjective Probabilities as a Reaction to Uncertainty'. Quarterly Journal of Economics, Vol. 75, pp. 670–689.
- Fellner, W. (1965). Probability and Profit: A Study of Economic Behavior along Bayesian Lines. Homewood, IL: Richard D. Irwin, Inc..
- Flavin, M. A. (1981). 'The Adjustment of Consumption to Changing Expectations About Future Income'. Journal of Political Economy, Vol. 89, pp. 975–1009.
- Fletcher, R. (1987). Practical Methods of Optimization. New York: John Wiley & Sons.
- Friedman, M. (1953). 'The Effects of a Full-Employment Policy on Economic Stability: A Formal Analysis'. In M. Friedman (ed.), Essays in Positive Economics University of Chicago Press, Chicago, IL.
- Friedman, M. (1956). A Theory of the Consumption Function. Princeton, NJ: Princeton University Press.
- Friedman, M. (1959). A Program for Monetary Stability. New York: Fordham University Press.

- Friedman, M. and L. J. Savage (1948). 'The Utility Analysis of Choices Involving Risk'. Journal of Political Economy, Vol. 56, pp. 279–304.
- Fudenberg, D. and D. M. Kreps (1995). 'Learning in Extensive Games, I: Self-Confirming and Nash Equilibrium'. Games and Economic Behavior, Vol. 8, pp. 20–55.
- Fudenberg, D. and D. M. Kreps (1995). 'Learning in Extensive Games, II: Experimentation and Nash Equilibrium'. Mimeo. Harvard University.
- Fudenberg, D. and D. K. Levine (1993). 'Self-Confirming Equilibrium'. Econometrica, Vol. 61, pp. 523–545.
- Fudenberg, D. and D. K. Levine (1995). 'Consistency and Cautious Fictitious Play'. Journal of Economic Dynamics and Control, Vol. 19, pp. 1065–1089.
- Fudenberg, D. and D. K. Levine (1998). The Theory of Learning in Games. Cambridge, MA: MIT Press.
- Gardiner, J. D. and A. J. Laub (1986). 'A Generalization of the Matrix-Sign-Function Solution for Algebraic Riccati Equations'. International Journal of Control, Vol. 44, pp. 823–832.
- Gardiner, J. D., M. R. Wette, A. J. Laub, J. J. Amato and C. B. Moler (1992). 'A FORTRAN-77 Software Package for Solving the Sylvester Matrix Equation $AXB^T + CXD^T = E$ '. ACM Transactions on Mathematical Software, Vol. 18, pp. 232–238.
- Giannoni, M. P. (2000). 'Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model'. Macroeconomic Dynamics, In press.
- Gilboa, I. and D. Schmeidler (1989). 'Maxmin Expected Utility with Non-Unique Prior'. Journal of Mathematical Economics, Vol. 18(2), pp. 141–153.
- Glover, K. and J. C. Doyle (1988). 'State-space Formulae for All Stabilizing Controllers that Satisfy an H_{∞} -Norm Bound and Relations to Risk-Sensitivity'. System & Control Letters, Vol. 11(1), pp. 167–172.
- Golub, G. H., S. Nash and C. Van Loan (1979). 'A Hessenberg-Schur Method for the Matrix Problem AX + XB = C'. *IEEE Transactions on Automatic Control*, Vol. AC-24, pp. 909–913.
- Golub, G. H. and C. Van Loan (1989). *Matrix Computations*. Baltimore, MD: Johns Hopkins University Press.
- Golub, G. H. and J. H. Wilkinson (1976). 'Ill-conditioned Eigensystems and the Computation of the Jordan Canonical Form'. SIAM Review, Vol. 18, pp. 578–619.
- Goodfriend, M. and R. King (1997). 'The New Neoclassical Synthesis and the Role of Monetary Policy'. In B. S. Bernanke and J. J. Rotemberg (eds.), NBER Macroeconomics Annual, 1997. The MIT Press, Cambridge, MA, pp. 231–282.
- Gudmundsson, T., C. Kenney and A. J. Laub (1992). 'Scaling of the Discrete-Time Algebraic Riccati Equation to Enhance Stability of the Schur Method'. *IEEE Transactions* on Automatic Control, Vol. 37, pp. 513–518.
- Hall, R. E. (1978). 'Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence'. *Journal of Political Economy*, Vol. 86, pp. 971–987.
- Hamilton, J. D. (1989). 'A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle'. Econometrica, Vol. 57, pp. 357–384.
- Hamilton, J. D. (1994). Time Series Analysis. Princeton, NJ: Princeton University Press.
- Hammarling, S. J. (1982). 'Numerical Solution of the Stable Nonnegative Lyapunov Equation'. IMA Journal of Numerical Analysis, Vol. 2, pp. 303–323.

References

- Hansen, L. P. (1987). 'Calculating Asset Prices in Three Exchange Economies'. Advances in Econometrics, Fifth World Congress, Cambridge, MA: Cambridge University Press.
- Hansen, L. P., D. Epple and W. Roberds (1985). 'Linear-Quadratic Duopoly Models of Resource Depletion'. In T. J. Sargent (ed.), *Energy, Foresight, and Strategy Resources* for the Future, Washington, D.C..
- Hansen, L. P., J. Heaton and T. J. Sargent (1991). 'Faster Methods for Solving Continuous Time Recursive Linear Models of Dynamic Economies'. Rational Expectations Econometrics, pp. 177–208.
- Hansen, L. P. and R. Jagannathan (1991). 'Implications of Security Market Data for Models of Dynamic Economies'. Journal of Political Economy, Vol. 99, pp. 225–262.
- Hansen, L. P. and T. J. Sargent (1993). 'Seasonality and Approximation Errors in Rational Expectations Models'. Journal of Econometrics, Vol. 55, pp. 21–55.
- Hansen, L. P. and T. J. Sargent (1994). 'Recursive Linear Models of Dynamic Economies'. Mimeo. University of Chicago.
- Hansen, L. P. and T. J. Sargent (1995). 'Discounted Linear Exponential Quadratic Gaussian Control'. IEEE Transactions on Automatic Control, Vol. 40, pp. 968–971.
- Hansen, L.P., T. J. Sargent, G. A. Turmuhambetova, and Noah Williams (2001). 'Robustness and Uncertainty Aversion'. Mimeo. University of Chicago.
- Hansen, L.P., and T.J. Sargent (2003). 'Robust control of forward-looking models'. Journal of Monetary Economics, Vol. ??, pp. ??.
- Hansen, L. P., T. J. Sargent and T. Tallarini (1999). 'Robust Permanent Income and Pricing'. Review of Economic Studies, Vol. 66, pp. 873–907.
- Hansen, L. P., T. J. Sargent and N. E. Wang (2000). 'Robust Permanent and Income and Pricing with Filtering'. *Macroeconomic Dynamics*, In press.
- Hansen, L.P., T. Sargent, G. Turmuhambetova, N. Williams (2001). 'Robustness and Uncertainty Aversion'. Mimeo. University of Chicago.
- Heaton, J. (1993). 'The Interaction Between Time-Nonseparable Preferences and Time Aggregation'. Econometrica, Vol. 61(2), pp. 353–385.
- Hitz, K. L. and B. D. O. Anderson (1972). 'Iterative Method of Computing the Limiting Solution of the Matrix Riccati Differential Equation'. Proc. of IEEE, Vol. 60, pp. 1402–1406.
- Jacobson, D. J. (1973). 'Optimal Linear Systems with Exponential Performance Criteria and their Relation to Differential Games'. *IEEE Transactions on Automatic Control*, Vol.18, pp. 124–131.
- James, Matthew R. (1992). 'Asymptotic Analysis of Nonlinear Stochastic Risk Sensitive Control and Differential Games'. journal=mathctrlsigsys, Vol. 5, pp. 401–417.
- Jovanovic, B. (1979). 'Job Matching and the Theory of Turnover'. Journal of Political Economy, Vol. 87(5), pp. 972–990.
- Jovanovic, B. and Y. Nyarko (1996). 'Learning by Doing and the Choice of Technology'. Econometrica, Vol. 64(6), pp. 1299–1310.
- Kågström, B. and P. Poromaa (1994). 'Computing Eigenspaces with Specified Eigenvalues of a Regular Matrix Pair (A, B) and Condition Estimation: Theory, Algorithms and Software'. Mimeo. LAPACK Working Note 87.
- Kasa, K. (1998). 'Model Uncertainty, Robust Policies, and the Value of Commitment'. Mimeo. Federal Reserve Bank of San Francisco.

- Kasa, K. (2000). 'Model Uncertainty, Robust Policies, and the Value of Commitment'. Macroeconomic Dynamics, In press.
- Kashyap, R. L. (1970). 'Maximum Likelihood Identification of Stochastic Linear Systems'. IEEE Transactions on Automatic Control, Vol. AC-15, pp. 25–34.
- Kenney, C. S., A. J. Laub and P. M. Papadopoulos (1993). 'A Newton-Squaring Algorithm for Computing the Negative Invariant subspace of a Matrix'. *IEEE Transactions on Automatic Control*, Vol. 38, pp. 1284–1289.
- Keynes, J. M. (1936). The General Theory of Employment, Interest, and Money. London: Macmillan.
- Kimball, M. S. (1990). 'Precautionary Saving in the Small and in the Large'. Econometrica, Vol. 58, pp. 53–73.
- Kimura, M. (1988). 'Convergence of the Doubling Algorithm for the Discrete-Time Algebraic Riccati Equation'. Int. J. Systems Sci., Vol. 19(5), pp. 701–711.
- Kimura, M. (1989). 'Doubling Algorithm for Continuous-Time Algebraic Riccati Equation'. Int. J. Systems Sci., Vol. 20(2), pp. 191–202.
- King, R. G., C. I. Plosser and S. T. Rebelo (1988a). 'Production, Growth, and Business Cycles: I. The Basic Neoclassical Model'. *Journal of Money, Credit, and Banking*, Vol. 21(2/3), pp. 195–232.
- King, R. G., C. I. Plosser and S. T. Rebelo (1988b). 'Production, Growth, and Business Cycles: II. New Directions'. Journal of Money, Credit, and Banking, Vol. 21(2/3), pp. 309–342.
- King, R. G. and A. L. Wolman (1999). 'What Should the Monetary Authority Do When Prices are Sticky?'. In J. B. Taylor (ed.), *Monetary Policy Rules Chicago: University* of Chicago Press, pp. 349–398.
- Knight, F. H. (1921). Risk, Uncertainty and Profit. Boston, MA: Houghton Mifflin Company.
- Kreps, D. M. (1998). 'Anticipated Utility and Dynamic Choice'. In D. P. Jacobs, E. Kalai and M. I. Kamien (eds.), Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983-1997. Cambridge University Press.
- Kwakernaak, H. and R. Sivan (1972). *Linear Optimal Control Systems*. New York: Wiley Interscience.
- Kydland, Finn E., and Edward C. Prescott. '1977'. Rules Rather than Discretion: The Inconsistency of Optimal Plans., Journal of Political Economy.Vol. 85(3), pp. 473–491
- Kydland, F. and E. C. Prescott (1980). 'Dynamic Optimal Taxation, Rational Expectations and Optimal Control'. Journal of Economic Dynamics and Control, Vol. 2(1), pp. 79–91.
- Kydland, F. and E. C. Prescott (1982). 'Time to Build and Aggregate Fluctuations'. Econometrica, Vol. 50, pp. 1345–1370.
- Laub, A. J. (1979). 'A Schur Method for Solving Algebraic Riccati Equations'. IEEE Trans. Auto. Control, Vol. AC-24, pp. 913–921.
- Laub, A. J. (1991). 'Invariant Subspace Methods for the Numerical Solution of Riccati Equations'. In S. Bittanti, A. J. Laub and J. C. Willems (eds.), *The Riccati Equation*. Springer Verlag, pp. 163–196
- Leland, H. E. (1968). 'Saving and Uncertainty: The Precautionary Demand for Saving'. Quarterly Journal of Economics, Vol. 82(3), pp. 465–473.

- Levin, A., V. Wieland and J. C. Williams (1999). 'Robustness of Simple Monetary Policy Rules under Model Uncertainty'. In J. B. Taylor (ed.), *Monetary Policy Rules Chicago:* University of Chicago Press, pp. 263–299.
- Lu, L. and W. Lin (1993). 'An Iterative Algorithm for Solution of the Discrete-Time Algebraic Riccati Equation'. *Linear Algebra and Its Applications*, Vol. 188,189, pp. 465–488.
- Lucas, R. E., Jr. (1976). 'Econometric Policy Evaluation: A Critique'. In K. Brunner and A. H. Meltzer (eds.), The Phillips Curve and Labor Markets. Amersterdam: North-Holland, pp. 19–46.
- Lucas, Robert E., Jr., and Nancy Stokey. '1983'. Optimal Monetary and Fiscal Policy in an Economy without Capital., Journal of Monetary Economics.Vol. 12(1), pp. 55–94
- Luenberger, David G. 1969. Optimization by Vector Space Methods.New York: John Wiley and Sons, Inc.
- MacFarlane, A. G. J. (1963). 'An Eigenvector Solution of the Optimal Linear Regulator Problem'. J. Electron. Contr., Vol. 14, pp. 643–654.
- Marcet, A. and R. Marimon (1992). 'Communication, Commitment, and Growth'. Journal of Economic Theory, Vol. 58(2), pp. 219–249.
- Marcet, A. and R. Marimon (1999). 'Recursive Contracts'. Mimeo. Universitat Pompeu Fabra, Barcelona.
- Marcet, A. and J. P. Nicolini (1998). 'Recurrent Hyperinflations and Learning'. Mimeo. Center for Economic Policy Research, Discussion Paper 1875.
- McCallum, B. T. (1997). 'Comment'. In B. S. Bernanke and J. J. Rotemberg (eds.), NBER Macroeconomics Annual, 1997. The MIT Press, Cambridge, MA, pp. 355–360.
- McCallum, B. T. and W. Nelson (1999). 'Performance of Operational Policy Rules in an Estimated Semiclassical Structural Model'. In J. B. Taylor (ed.), *Monetary Policy Rules Chicago: University of Chicago Press, pp. 15–45.*
- McGrattan, E. (1994). 'A Note on Computing Competitive Equilibria in Linear Models'. Journal of Economic Dynamics and Control, Vol. 18, pp. 149–160.
- McGrattan, E., R. Rogerson and R. Wright (1995). 'An Equilibrium Model of the Business Cycle with Household Production and Fiscal Policy'. Mimeo. Staff Report 166, Federal Reserve Bank of Minneapolis.
- Miller, B. L. (1974). 'Optimal Consumption with a Stochastic Income Stream'. Econometrica, Vol. 42(2), pp. 253–266.
- Miller, M. and M. Salmon (1985a). 'Dynamic Games and the Time Inconsistency of Optimal Policy in Open Economies'. *Economic Journal, Supplement*, Vol. 95(0), pp. 124–137.
- Miller, M. and M. Salmon (1985b). 'Policy Coordination and Dynamic Games'. In W. Buiter and R. Marston (eds.), *International Economic Policy Coordination*. Cambridge University Press, pp. 184–213.
- Mustafa, D. and K. Glover (1990). Minimum Entropy H^∞ Control. Berlin: Springer-Verlag.
- Muth, J. F. (1960). 'Optimal Properties of Exponentially Weighted Forecasts'. Journal of the American Statistical Association, Vol. 55(290), pp. 299–306.
- Muth, J. F. (1961). 'Rational Expectations and the Theory of Price Movements'. Econometrica, Vol. 29, pp. 315–335.

- Onatski, A. and J. H. Stock (2000). 'Robust Monetary Policy Under Model Uncertainty in a Small Model of the U.S. Economy'. *Macroeconomic Dynamics*, In press.
- Otrok, C. (In press). 'On measuring the welfare cost of business cycles'. Journal of Monetary Economics, In press.
- Otrok, C. (In press). 'Spectral Welfare Cost Functions'. International Economic Review, In press.
- Pappas, T., A. J. Laub and N. R. Sandell, Jr. (1980). 'On the Numerical Solution of the Discrete-Time Algebraic Riccati Equation'. *IEEE Transactions on Automatic Control*, Vol. AC-25(4), pp. 631–641.
- Pearlman, J. G. (1992). 'Reputational and Nonreputational Policies Under Partial Information'. Journal of Economic Dynamics and Control, Vol. 16(2), pp. 339–358.
- Pearlman, J. G., D. A. Currie and P. L. Levine (1986). 'Rational Expectations Models with Partial Information'. *Economic Modeling*, Vol. 3(2), pp. 90–105.
- Petkov, P. Jr., N. D. Christov and M. M. Konstantinov (1991). Computational Methods for Linear Control Systems. Englewood Cliffs, NJ: Prentice Hall.
- Potter, J. E. (1966). 'Matrix Quadratic Solutions'. SIAM Journal on Applied Mathematics, Vol. 14, pp. 496–501.
- Rabin, M. (1999). 'Risk Aversion and Expected-Utility Theory: A Calibration Theorem'. Mimeo. January, ????? .
- Roberts, J. D. (1980). 'Linear Model Reduction and Solution of the Algebraic Equation by Use of the Sign Function'. International Journal of Control, Vol. 32, pp. 677-687. (reprint of Technical Report No. TR-13, CUED/B-Control, Cambridge University, Engineering Department, 1971)
- Rosen, S., K. M. Murphy and J. A. Scheinkman (1994). 'Cattle Cycles'. Journal of Political Economy, Vol. 102(3), pp. 468–492.
- Rotemberg, J. and M. Woodford (1997). 'An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy'. In B. S. Bernanke and J. J. Rotemberg (eds.), *NBER Macroeconomics Annual*, 1997. The MIT Press, Cambridge, MA, pp. 297–346.
- Rotemberg, J. and M. Woodford (1999). 'Interest Rate Rules in an Estimated Sticky Price Model'. In J. B. Taylor (ed.), *Monetary Policy Rules* Chicago: University of Chicago Press, pp. 57–119.
- Rozanov, Y. A. (1967). Stationary Random Processes. San Francisco: Holden-Day.
- Ryoo, Jaewoo and Sherwin Rosen (2003). 'The Engineering Labor Market'. Mimeo. Hong Kong University of Science and Technology.
- Salmon, M. (1996). 'On the Theoretical Irrelevance of the Lucas Critique Twenty Years On'. Mimeo. European University Institute.
- Sargent, T. J. (1981). 'Interpreting Economic Time Series'. Journal of Political Economy, Vol. 89(2), pp. 213–248.
- Sargent, T. J. (1987). Macroeconomic Theory. 2nd ed. New York: Academic Press.
- Sargent, T. J. (1999). The Conquest of American Inflation. Princeton, NJ: Princeton University Press.
- Sargent, T. J. (1999). 'Comment'. In J. B. Taylor (ed.), Monetary Policy Rules Chicago: University of Chicago Press, pp. 144–154.
- Sargent, T. J. and N. Wallace (1975). 'Rational Expectations, the Optimal Monetary Instrument, and the Optimal Supply of Money'. *Journal of Political Economy*, Vol. 83, pp. 241–254.

- Sargent, T. J. and N. Wallace (1976). 'Rational Expectations and the Theory of Economic Policy'. Journal of Monetary Economics, Vol. 2(2), pp. 169–183.
- Savage, L. J. (1954). The Foundations of Statistics. New York: John Wiley & Sons.
- Segal, U. and A. Spivak (1990). 'First Order Versus Second Order Risk Aversion'. Journal of Economic Theory, Vol. 51, pp. 111–125.
- Shafer, G. (1976). A Mathematical Theory of Evidence. Princeton, NJ: Princeton University Press.
- Sims, C. A. (1972). 'Approximate Prior Restrictions in Distributed Lag Estimation'. Journal of the American Statistical Association, Vol. 67, pp. 169–175.
- Sims, C. A. (1974). 'Seasonality in Regression'. Journal of the American Statistical Association, Vol. 69, pp. 618–626.
- Sims, C. A. (1980). 'Macroeconomics and Reality'. Econometrica, Vol. 48(1), pp. 1–48.
- Sims, C. A. (1993). 'Rational Expectations Modeling with Seasonally Adjusted Data'. Journal of Econometrics, Vol. 55, pp. 9–19.
- Siow, A. (1984). 'Occupational Choice Under Uncertainty'. Econometrica, Vol. 52(3), pp. 631–645.
- Stewart, G. W. (1972). 'On the Sensitivity of the Eigenvalue Problem $Ax = \lambda Bx$ '. SIAM Journal on Numerical Analysis, Vol. 9, pp. 669–668.
- Stewart, G. W. (1976). 'Algorithm 506 HQR3 and EXCHNG: FORTRAN Subroutines for Calculating and Ordering the Eigenvalues of a Real Upper Hessenberg Matrix'. ACM Transactions on Mathematical Software, Vol. 2, pp. 275–280.
- Stock, J. H. (1999). 'Comment'. In J. B. Taylor (ed.), Monetary Policy Rules Chicago: University of Chicago Press, pp. 253–259.
- Stokey, Nancy L. '1989'. Reputation and Time Consistency., American Economic Review.Vol. 79, pp. 134–139
- Stokey, Nancy L. '1991'. Credible Public Policy., Journal of Economic Dynamics and Control.Vol. 15(4), pp. 627–656
- Svensson, L. E. O. and M. Woodford (2000). 'Indicator Variables for Monetary Policy'. Mimeo. Princeton University.
- Tallarini, T. D., Jr. (2000). 'Risk-Sensitive Real Business Cycles'. Journal of Monetary Economics, Vol. 45(3), pp. 507–532.
- Taylor, J. B. (1999). Monetary Policy Rules. Chicago, IL: University of Chicago Press.
- Tornell, A. (1998). 'Excess volatility of asset prices with H_{∞} forecasts'. Mimeo. Cambridge, MA, Harvard University.
- United States, Bureau of the Census (1975). Historical Statistics of the United States, Colonial Times to 1970. Washington, D.C.: U.S. Department of Commerce.
- United States, Bureau of the Census (1989). Agricultural Statistics. Washington, D.C.: U.S. Department of Commerce.
- Van Dooren, P. (1981). 'A Generalized Eigenvalue Approach for Solving Riccati Equations'. SIAM Journal on Scientific and Statistical Computing, Vol. 2, pp. 121–135.
- Van Dooren, P. (1982). 'Algorithm 590-DSUBSP and EXCHQZ: Fortran Subroutines for Computing Deflating Subspaces with Specified Spectrum'. ACM Transactions on Mathematical Software, Vol. 8, pp. 376–382.
- Vaughan, D. R. (1970). 'A Nonrecursive Algebraic Solution for the Discrete Riccati Equation'. IEEE Transactions on Automatic Control, Vol. AC-15(5), pp. 597–599.

Vuong, Q.H. (1989). 'Likelihood ratio tests for model selection and non-nested hypotheses'. Econometrica, Vol. 57, No. 2,March, pp. 307-333.

- von Neumann, J. and O. Morgenstern (1944). Theory of Games and Economic Behavior. Princeton, NJ: Princeton University Press.
- von zur Muehlen, P. (1982). 'Activist Vs. Non-Activist Monetary Policy: Optimal Rules Under Extreme Uncertainty'. Mimeo. Board of Governors of the Federal Reserve Board, Washington, D.C..
- Wallis, K. F. (1974). 'Seasonal Adjustment and Relations Between Variables'. Journal of the American Statistical Association, Vol. 69, pp. 18–31.
- Wang, T. (1999). 'Updating Rules for Non-Bayesian Preferences'. Mimeo. ????? .
- Ward, R. C. (1981). 'Balancing the Generalized Eigenvalue Problem'. SIAM Journal on Scientific and Statistical Computing, Vol. 2, pp. 141–152.
- Weil, P. (1989). 'The Equity Premium Puzzle and the Risk-Free Rate Puzzle'. Journal of Monetary Economics, Vol. 24(2), pp. 401–421.
- Weil, P. (1990). 'Nonexpected Utility in Macroeconomics'. Quarterly Journal of Economics, Vol. 105, pp. 29–42.
- Weil, P. (1993). 'Precautionary Savings and the Permanent Income Hypothesis'. Review of Economic Studies, Vol. 60, pp. 367–383.
- Whiteman, C. (1983). Linear Rational Expectations Models: A User's Guide. Minneapolis, MN: University of Minnesota Press.
- Whiteman, C. (1985a). 'Analytical Policy Design Under Rational Expectations'. Econometrica, Vol. 54(6), pp. 1387–1405.
- Whiteman, C. (1985b). 'Spectral Utility, Wiener-Hopf Techniques, and Rational Expectations'. Journal of Economic Dynamics and Control, Vol. 9(2), pp. 225–240.
- Whiteman, C. (1986). 'Analytical Policy Design Under Rational Expectations'. Econometrica, Vol. 54, pp. 1387–1405.
- Whittle, P. (1983). Prediction and Regulation by Linear Least-Square Methods, Second edition, revised. Minneapolis, MN: University of Minnesota Press.
- Whittle, P (1990). Risk-Sensitve Optimal Control. New York: Wiley.
- Whittle, P. (1996). Optimal Control: Basics and Beyond. New York: Wiley.
- Wieland, V. (1996). 'Monetary policy, Parameter Uncertainty, and Optimal Learning'. Mimeo. Board of Governors of the Federal Reserve System.
- Wilson, D. A. and A. Kumar (1982). 'Derivative Computations for the Log Likelihood Function'. *IEEE Transactions on Automatic Control*, Vol. AC-27, pp. 230–232.
- Woodford, M. (1998). 'Optimal Monetary Policy Inertia'. Mimeo. Princeton University, December 14.
- Zadrozny, P. A. (1988a). 'Analytic Derivatives for Estimation of Discrete-Time, Linear Quadratic, Dynamic, Optimization Models'. *Econometrica*, Vol. 56, pp. 467–472.
- Zadrozny, P. A. (1988b). 'Gaussian Likelihood of Continuous-Time ARMAX Models when Data are Stocks and Flows at Different Frequencies'. *Econometric Theory*, Vol. 4, pp. 108–124.
- Zadrozny, P. A. (1989). 'Analytic Derivatives for Estimation of Linear Dynamic Models'. Computers and Mathematics with Applications, Vol. 18, pp. 539–553.
- Zadrozny, P. A. (1992). 'Errata to Analytic Derivatives for Estimation of Linear Dynamic Models'. Computers and Mathematics with Applications, Vol. 24, pp. 289–290.

- Zeldes, S. P. (1989). 'Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence'. Quarterly Journal of Economics, Vol. 104(2), pp. 275–298.
- Zhou, K., J. C. Doyle and K. Glover (1996). Robust and Optimal Control. London: Prentice Hall.

Index

absolute continuity, 135 backward shift operator, 311 Bayesian interpretation and belief heterogeneity, 363 Bellman equation Kalman filter, 109 optimal linear regulator, 20 robust optimal linear regulator, 26 - 39Bergemann, Dirk, 126 canonical representation of preferences, 307 certainty equivalence, 25, 145, 213, 222 for problem without robustness, 21modified for robustness, 27 concavity and entropy criterion, 197 constraint game dynamic, 170 pathology, 130 Stackelberg, 370 static version, 120-136 continuous time, 88 controllablity, 61 \mathcal{D} , 39, 163, 169, 196, 393–394 deflating subspace, 73 stable, 159 detectability, 62 detection error probability, 213-219

distortion in covariance, 401 in mean, 401 doubling algorithm Riccati equation, 77, 82-88 Sylvester equations, 96 duality, 114 Kalman filter and linear regulator, 99, 316 robust control and filter, 319 dynamic programming, 17 economy information, 251-266, 294 preferences, 251-266, 294 technology, 251-266, 294 entropy, 22-24, 112, 113 criterion, 180, 197 measuring model misspecification, 7 misspecification analysis in econometrics, 7 entropy criterion, 198 etymology, 194 equilibrium, 251 competitive, 259 rational expectations, 306 equity premium, 301 extremization, 28 forecasting problem, 172, 231 Fourier transform, 180 golden ratio, 328 H_2 , 180, 181, 183, 197–198, 210 H_{∞} , 135, 180, 183–184, 198 Hamiltonian matrix, 75 Hessenberg-Schur algorithm, 94

- 439 -

Index

heterogeneity of worst case models, 378 implementability multipliers, 50 initial condition, 190 $H_2, 183$ $H_{\infty}, 184$ Stackelberg constraint problem, 180invariant subspace, 70 Kalman filter, 100-115 Lagrange multiplier theorem, 126, 141, 163 market price of risk, 301 Markov perfect equilibrium concerns about robustness, 361 matrix sign algorithm, 77, 88-90 misspecification in Ramsey problem, 371 multiplier game in sequences, 146 Markov perfect, 161–163 Stackelberg, 161, 371 static version, 120-136 observability, 63 occupational choice, 305 optimal linear regulator problem, 20, 60, 61, 64, 65, 99, 171, 182, 256, 336 augmented, 60, 63-64, 90, 107, 108, 159 Parseval's equality, 181 pencil, 72 permanent income model, 40, 47, 221 - 248

Phillips curve, 119 policy improvement algorithm, 104, 142precautionary savings, 13, 42-47, 222, 233 - 244principle components, 183 \mathcal{R} , 195, 394–402 Radon-Nikodym derivative, 293, 309, 407Ramsey problem, 367 first-order approach, 412 optimal taxation, 411 robust, 368, 371, 416 rational expectations, 137 rational expectations econometrics, 7 relative entropy, 396-402 Riccati equation, 69–75, 158, 170 Kalman filter, 110 risk sensitivity, 34, 178, 195, 393-402 robustness bound frequency domain, 187 nonlinear version, 399 time domain, 32, 171 Schur algorithm, 77–82 shadow price, 47, 372 spectral density factorization, 188, 191 robustness flattens, 332 spectral density matrix factorization, 191 stabilizability, 61 Stackelberg game, 367, 368 frequency domain, 182

multiplier, 186

time domain, 179 stochastic discount factor, 297 subspace deflating, 72 Sylvester equation, 90, 92, 93, 96, 257, 269, 282, 297, 299, 411 efficient solution, 93–94 symplectic matrix, 70 pencil, 148, 373, 389 T, 39, 163, 170, 393 T, 395–402 taxation, 406–425 time-inconsistency and conflict, 377

Author Index

Anderson, B.D.O., 47, 48, 69, 81 Anderson, Evan, 47, 183

Başar, T., 294, 318
Ball, L., 187
Bernhard, Pierre, 294
Brock, William, 235

Calvino, Italo, 12, 82 Christiano, Lawrence, 318 Currie, David, 318

Doyle, John C., 1, 164, 166

Epple, Dennis, 318, 320 Evans, George, 2

Friedman, Milton, 83

Giannoni, Mark, 318 Glover, Keith, 1, 164, 166 Golub, G.H., 47

Hansen, Lars Peter, 4, 9, 183, 223, 318, 320 Honkapohja, Seppo, 2

Jacobson, D.J., 9, 21, 28, 32, 102, 168 Jevons, William Stanley, 293 Jones, Larry, 365

Kasa, Kenneth, 318 King, Robert G., 318 Kydland, Finn, 98

Laub, A.J., 47 Levine, Paul, 318 Lucas, Robert E., Jr., 4, 351 Luenberger, David, 104, 117, 139 Manuelli, Rodolfo, 365 Marcet, Albert, 2, 318 Marimon, Ramon, 318 McGrattan, Ellen R., 47 Mehra, Raj, 235 Miller, Marcus, 318 Moore, J.B., 48, 69, 81 Muth, John F., 83 Nash, S., 47 Pappas, T., 47 Pearlman, J.G., 318, 320 Prescott, Edward C., 98, 235 Roberds, William, 318, 320 Rosen, Sherwin, 267 Rossi, Peter, 365 Rozanov, Yu, 162, 164 Ryuoo, Jaewoo, 267 Salmon, Mark, 318 Sandell, N.R., 47 Sargent, Thomas J., 4, 9, 183, 187, 223 Siow, Aloyisius, 267 Stokey, Nancy, 99, 114, 351 Van Loan, C., 47 Velde, François, 351 Whittle, Peter, 9, 21, 32, 82, 87, 88, 102, 163, 294 Woodford, Michael, 2, 318

Zhou, Kemin, 1, 164, 166

- 443 -

Matlab Index