

---

# ROBUSTNESS



---

# ROBUSTNESS

LARS PETER HANSEN

THOMAS J. SARGENT

PRINCETON UNIVERSITY PRESS    PRINCETON AND OXFORD

© 2008 by Princeton University Press  
Published by Princeton University Press,  
41 William Street, Princeton, New Jersey 08540

In the United Kingdom: Princeton University Press  
3 Market Place, Woodstock, Oxfordshire, OX20 1SY

All Rights Reserved

Library of Congress Control Number: 2007927641  
ISBN-13: 978-0-691-11442-2 (cloth)

British Library Cataloguing-in-Publication Data are available

The authors composed this book in Computer Modern  
using  $\text{\TeX}$  and the  $\text{\TeX}$ sis 2.18 macros.

Printed on acid-free paper.  $\infty$

[press.princeton.edu](http://press.princeton.edu)

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

In memory of our friend Sherwin Rosen



# Contents

Preface vii

Acknowledgments ix

## Part I: Motivation and main ideas

1. Introduction 3

Generations of control theory. Control theory and rational expectations. Misspecification and rational expectations. Our extensions of robust control theory. Discounting. Representation of worst-case shock. Multiple agent settings. Explicitly stochastic interpretations. Calibrating fear of misspecification. Robust filtering and estimation. Robust control theory, shock serial correlations, and rational expectations. Entropy in specification analysis. Acknowledging misspecification. Why entropy? Why max-min? Is max-min too cautious? Aren't you just picking a plausible prior? Why not learn the correct specification? Is the set of perturbed models too limited? Is robust control theory positive or normative? Other lessons. Topics and organization.

2. Basic ideas and methods 25

Introduction. Approximating models. Dynamic programming without model misspecification. Measuring model misspecification with entropy. Two robust control problems. Modified certainty equivalence principle. Robust linear regulator. More general misspecifications. A simple algorithm. Interpretation of the simple algorithm. Robustness and discounting in a permanent income model. The LQ permanent income model. Solution when  $\sigma = 0$ . Linear regulator for permanent income model. Effects on consumption of concern about misspecification. Observational equivalence of quantities but not continuation values. Distorted endowment process. A Stackelberg formulation for representing misspecification. Concluding remarks. Matlab programs.

### 3. A stochastic formulation 53

Introduction. Shock distributions. Martingale representations of distortions. A digression on entropy. A stochastic robust control problem. A recursive formulation. Verifying the solution. A value function bound. Large deviation interpretation of  $\mathcal{R}$ . Choosing the control law. Linear-quadratic model. Relative entropy and normal distributions. Value function adjustment for the LQ model.

## Part II: Standard control and filtering

### 4. Linear control theory 67

Introduction. Control problems. Deterministic regulator problem. Augmented regulator problem. Discounted stochastic regulator problem. Solving the deterministic linear regulator problem. Nonsingular  $A_{yy}$ . Singular  $A_{yy}$ . Continuous-time systems. Computational techniques for solving Riccati equations. Schur algorithm. Digression: solving DGE models with distortions. Doubling algorithm. Initialization from a positive definite matrix. Application to continuous time. Matrix sign algorithm. Solving the augmented regulator problem. Computational techniques for solving Sylvester equations. The Hessenberg-Schur algorithm. Doubling algorithm. Conclusion.

### 5. The Kalman filter 103

Introduction. Review of Kalman filter and preview of main result. Muth's problem. The dual to Muth's filtering problem. The filtering problem. Sequence version of primal and dual problems. Sequence version of Kalman filtering problem. Sequence version of dual problem. Digression: reversing the direction of time. Recursive version of dual problem. Recursive version of Kalman filtering problem. Concluding remarks.

## Part III: Robust control

### 6. Static multiplier and constraint games 119

Introduction. Phillips curve example. The government's problem. Robustness of robust decisions. Basic setup with a correct model. The constraint game with  $b = 0$ . Multiplier game with  $b = 0$ . The model with  $b \neq 0$ . Probabilistic formulation ( $b = 0$ ). Gaussian perturbations. Letting the minimizing agent choose random perturbations



when  $b = 0$ . Constraint and multiplier preferences. Concluding remarks. Rational expectations equilibrium.

## 7. Time domain games for attaining robustness 139

Alternative time domain formulations. The setting. Two Stackelberg games. Two Markov perfect equilibria. Markov perfect equilibrium: definition. Markov perfect equilibria: value functions. Useful recursions. Computing a Markov perfect equilibrium: the recursion. Step one: distorting the covariance matrix. Step two: distorting the mean. Another Markov perfect equilibrium and a Bellman-Isaacs condition. Taking inventory. Markov perfect equilibrium of infinite horizon game. Recursive representations of Stackelberg games. Markov perfect equilibria as candidate Stackelberg equilibria. Maximizing player chooses first. Minimizing player chooses first. Bayesian interpretation of robust decision rule. Relation between multiplier and constraint Stackelberg problems. Dynamic consistency. Miscellaneous details. Checking for the breakdown point. Policy improvement algorithm. Concluding remarks. Details of a proof of Theorem 7.7.1. Certainty equivalence. Useful formulas. A single Riccati equation. Robustness bound. A pure forecasting problem. Completing the square.

## 8. Frequency domain games and criteria for robustness 173

Robustness in the frequency domain. Stackelberg game in time domain. Fourier transforms. Stackelberg constraint game in frequency domain. Version 1:  $H_2$  criterion. Version 2: the  $H_\infty$  criterion. Stackelberg multiplier game in frequency domain. A multiplier problem. Robustness bound. Breakdown point reconsidered. Computing the worst-case misspecification. Three examples of frequency response smoothing. Example 1. Example 2. Example 3. Entropy is the indirect utility function of the multiplier game. Meaning of entropy. Risk aversion across frequencies. Concluding remarks. Infimization of  $H_\infty$ . A dual prediction problem. Proofs of three lemmas. Duality. Evaluating a given control law. When  $\theta = \theta_F$ . Failure of entropy condition. Proof of Theorem 8.8.2. Stochastic interpretation of  $H_2$ . Stochastic counterpart.

## 9. Calibrating misspecification fears with detection error probabilities 213

Introduction. Entropy and detection error probabilities. The context-specific nature of  $\theta$ . Approximating and distorting models. Detection error probabilities. Details. Likelihood ratio under the approximating model. Likelihood ratio under the distorted model. The detection error probability. Breakdown point examples revisited. Ball's model. Concluding remarks.

## 10. A permanent income model 223

Introduction. A robust permanent income theory. Solution when  $\sigma = 0$ . The  $\sigma = 0$  benchmark case. Observational equivalence for quantities of  $\sigma = 0$  and  $\sigma \neq 0$ . Observational equivalence: intuition. Observational equivalence: formal argument. Precautionary savings interpretation. Observational equivalence and distorted expectations. Distorted endowment process. Another view of precautionary savings. Frequency domain representation. Detection error probabilities. Robustness of decision rules. Concluding remarks. Parameter values. Another observational equivalence result.

## Part IV: Multi-agent problems

## 11. Competitive equilibria without robustness 253

Introduction. Pricing risky claims. Types of competitive equilibria. Information, preferences, and technology. Information. Preferences. Technology. Planning problem. Imposing stability. Arrow-Debreu equilibrium. The price system at time 0. The household. The firm. Competitive equilibrium with time-0 trading. Equilibrium computation. Shadow prices. Recursive representation of time 0 prices. Recursive representation of household's problem. Units of prices and reopening markets. Sequential markets with Arrow securities. Arrow securities. The household's problem in the sequential equilibrium. The firm. Recursive competitive equilibrium. Asset pricing in a nutshell. Partial equilibrium interpretation. Concluding remarks.

## 12. Competitive equilibria with robustness 271

Introduction. A pure endowment economy. The planning problem. Household problem. A robust planning problem. Max-min representation of household problem. Sequence problem of maximizing player. Digression about computing  $\mu_0^w$ . Sequence problem of minimizing player. A decentralization with Arrow securities. A robust consumer trading Arrow securities. The inner problem. The outer problem. An *ex post* Bayesian planning problem. Practical remarks. A model of occupational choice and pay. A one-occupation model. Equilibrium with no concern about robustness. Numerical example of Ryoo-Rosen model. Two asset pricing strategies. Pricing from the robust planning problem. Pricing from the *ex post* Bayesian planning problem. Concluding remarks. Decentralization of partial equilibrium. Solving Ryoo and Rosen's model by hand.

### 13. Asset pricing 295

Introduction. Approximating and distorted models. Asset pricing without robustness. Asset pricing with robustness. Adjustment of stochastic discount factor for fear of model misspecification. Reopening markets. Pricing single-period payoffs. Calibrated market prices of model uncertainty. Concluding remarks.

### 14. Risk sensitivity, model uncertainty, and asset pricing 307

Introduction. Organization. The equity premium and risk-free rate puzzles. Shocks and consumption plans. Recursive preferences. Stochastic discount factor for risk-sensitive preferences. Risk-sensitive preferences let Tallarini attain the Hansen-Jagannathan bounds. Reinterpretation of the utility recursion. Using martingales to represent probability distortions. Recursive representations of distortions. Ambiguity averse multiplier preferences. Observational equivalence? Yes and no. Worst-case random walk and trend stationary models. Market prices of risk and model uncertainty. Calibrating  $\gamma$  using detection error probabilities. Recasting Tallarini's graph. Concluding remarks. Value function and worst-case process. The value function. The distortion. An alternative computation. The trend stationary model.

### 15. Markov perfect equilibria with robustness 327

Introduction. Markov perfect equilibria with robustness. Explanation of  $x_{it+1}, x_t$  notation. Computational algorithm: iterating on stacked Bellman equations. Bayesian interpretation and belief heterogeneity. Heterogeneous worst-case beliefs. Concluding remarks.

### 16. Robustness in forward-looking models 333

Introduction. Related literature. The robust Stackelberg problem. Multiplier version of the robust Stackelberg problem. Solving the robust Stackelberg problem. Step 1: Solve a robust linear regulator. Step 2: Use the stabilizing properties of shadow price  $Py_t$ . Step 3: Convert implementation multipliers into state variables. Law of motion under robust Ramsey plan. A monopolist with a competitive fringe. The approximating and distorted models. The problem of a firm in the competitive fringe. Changes of measure. Euler equation for  $\lambda^q$  under the approximating model. Euler equation for  $\lambda^q$  under the monopolist's perturbed model. The monopolist's transition equations. The monopolist's problem. Computing the volatility loading on  $\lambda_t^q$ . Timing subtlety. An iterative algorithm. Recursive representation of a competitive firm's problem. Multipliers. Cross-checking the solution for  $u_t, w_{t+1}$ . Numerical example. Concluding remarks. Invariant subspace method. The Riccati equation. Another Bellman equation.

## Part V: Robust estimation and filtering

### 17. Robust filtering with commitment 359

Alternative formulations. A linear regulator. A static robust estimation problem. A digression on distorted conditional expectations. A dynamic robust estimation problem. Many-period filtering problem. Duality of robust filtering and control. Matlab programs. The worst-case model. Law of motion for the state reconstruction error. The worst-case model associated with a time-invariant  $K$ . A deterministic control problem for the worst-case mean. A Bayesian interpretation. Robustifying a problem of Muth. Reconstructing the ordinary Kalman filter. Illustrations. Another example. A forward looking perspective. Relation to a formulation from control theory literature. The next chapter. Dual to evil agent's problem.

### 18. Robust filtering without commitment 383

Introduction. A recursive control and filtering problem. The decision maker's approximating model. Two sources of statistical perturbation. Two operators. The  $\mathbf{T}^1$  operator. The  $\mathbf{T}^2$  operator. Two sources of fragility. A recursive formulation for control and estimation. A certainty equivalent shortcut. Computing the  $\mathbf{T}^1$  operator. Worst-case distribution for  $z - \tilde{z}$  is  $\mathcal{N}(u, \Gamma(\Delta))$ . Worst-case signal distribution. Examples. Concluding remarks. Worst-case signal distribution.

## Part VI: Extensions

### 19. Alternative approaches 403

Introduction. More structured model uncertainty. Model averaging. A shortcut. Probabilistic sophistication. Time inconsistency. Continuation entropy. Disarming the entropy constraint. Rectangularity can be taken too far.

### References 413

### Index 427

Author Index	431
Matlab Index	435



## Preface

A good decision rule for us has been, “if Peter Whittle wrote it, read it.” Whittle’s book, *Prediction and Regulation by Linear Least Squares Methods* (originally published in 1963, revised and reprinted in 1983), taught early builders and users of rational expectations econometrics, including us, the classical time series techniques that are 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 through them. These and other books on robust control theory, such as Başar and Bernhard’s 1995  *$H^\infty$  – Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*, provide tools for approaching the ‘soft’ but important question of how to make decisions when you don’t fully trust your model.

Work on robust control theory opens up the possibility of rigorously analyzing how agents should cope with fear of model 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 designed mainly for types of problems that differ significantly from economic problems. Therefore, we soon recognized that we would have to modify and extend aspects of risk-sensitive and robust control methods if we were to apply them to economic problems. That is why we started the research that underlies this book. We do not claim to have attained a general theory of how to make economic decisions in the face of model misspecification, but only to have begun to study this difficult and important problem that has concerned every researcher who has estimated and tried to validate a rational expectations model, every central banker who has knowingly used dubious models to guide his monetary policy decisions, and every macroeconomist whose specification doubts have made him regard formal estimation as wrongheaded and who has instead “calibrated” the parameters of a complete, but admittedly highly stylized, model.





## Acknowledgments

For criticisms of previous drafts and stimulating discussions of many issues we thank Fernando Alvarez, Francisco Barillas, Marco Bassetto, Luca Benati, Dirk Bergemann, V.V. Chari, Eugene Chiu, Richard Dennis, Jack Y. Favilukis, Anastasios Karantounias, Kenneth Kasa, Patrick Kehoe, Junghoon Lee, Francesco Lippi, Pascal Maenhout, Ricardo Mayer, Anna Orlik, Joseph Pearlman, Tomasz Piskorski, Mark Salmon, Christopher Sims, Jose Scheinkman, Martin Schneider, Tomasz Strzalecki, Joseph Teicher, Aaron Tornell, François Velde, Peter von zur Muehlen, Neng Wang, Yong Wang, Pierre-Oliver Weil, and Noah Williams. We especially thank Anna Orlik for reading and criticizing the entire manuscript. We also owe a special thanks to François Velde for extraordinary help with typesetting and design problems. We thank Evan Anderson and Ellen McGrattan for allowing us to use many of the ideas in our joint paper in chapter 4. In addition to providing comments, Francisco Barillas, Christian Matthes, Ricardo Mayer, Tomasz Piskorski, Yong Shin, Stijn Van Nieuwerburgh, Chao D. Wei, and Mark Wright helped with the computations. We thank the National Science foundation for separate grants that have supported our research. Sargent thanks William Berkley for several useful conversations about risk and uncertainty. Our editors at Princeton University Press, Dale Cotton, Seth Ditchik, and Peter Dougherty provided encouragement and valuable suggestions about style and presentation. We thank John Doyle, an artist as well as a manufacturer of robust control theory, for letting us reproduce figure 1.1.1. His menacing image of a robust control theorist brandishing  $\theta$  reminds us why Arthur Goldberger and Robert E. Lucas, Jr., warned us to beware of theorists bearing free parameters. Our aim is to convince readers that the parameter  $\theta$  provides a practical way to confront concerns about model misspecification that applied economists encounter daily. It is tempting to make the agents in our models fear misspecification too and to study the outcomes it induces. This book is about how to do that.



*Part I*

*Motivation and main ideas*



# Chapter 1

## Introduction

*Knowledge would be fatal, it is the uncertainty that charms one. A mist makes things beautiful.*

— Oscar Wilde, *The Picture of Dorian Gray*, 1891

### 1.1. Generations of control theory


Figure 1.1.1 reproduces John Doyle’s cartoon about developments in optimal control theory since World War II.<sup>1</sup> 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 filters). The scientists in the top panels completely trust their models of the transition dynamics. The, shall we say, gentleman in the lower panel shares the objectives of his predecessors from the 50s, 60s, and 70s, but regards his model 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 nondenumerable set of models near his approximating model. The  $H_\infty$  in his postmodern tattoo and the  $\theta$  on his staff are alternative ways to express doubts about his approximating model by measuring the discrepancy of the true data generating mechanism from his approximating model. As we shall learn in later chapters, the parameter  $\theta$  is interpretable as a penalty on a measure of discrepancy (entropy) between his approximating model and the model that actually generates the data. The  $H_\infty$  refers to the limit of his objective function as the penalty parameter  $\theta$  approaches a “break down point” that bounds the set of alternative models against which the decision maker can attain a robust decision rule.

### 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. A significant reason that rational expectations initially diffused slowly after Muth’s (1961) paper is that in 1961 few economists knew the tools lampooned in the top panel of figure 1.1.1. Rational expectations took

---

<sup>1</sup> John Doyle consented to let us reproduce this drawing, which appears in Zhou, Doyle, and Glover (1996). We changed Doyle’s notation by making  $\theta$  (Doyle’s  $\mu$ ) the free parameter carried by the post-modern control theorist.



**Figure 1.1.1:** A pictorial history of control theory (courtesy of John Doyle). Beware of a theorist bearing a free parameter,  $\theta$ .

hold in the 1970s only after a new generation of macroeconomists had learned those tools. Ever since, macroeconomists and rational expectations econometricians have gathered inspiration and ideas from classical and recursive control theory.<sup>2</sup>

When macroeconomists were beginning to apply classical and modern control and estimation theory in the late 1970s, control theorists and applied mathematicians were seeking ways to relax the assumption that the decision maker trusts his model. They sought new control and estimation methods to improve adverse outcomes that came 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 had promised and therefore sought decision rules and estimators that acknowledged model misspecification. That is how robust control and estimation theory came to be.

<sup>2</sup> See Stokey and Lucas with Prescott (1989), Ljungqvist and Sargent (2004), and Hansen and Sargent (1991) for many examples.

### 1.3. Misspecification and rational expectations

To say that model misspecification is as much of a problem in economics as it is in physics and engineering is an understatement. 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. We assume that a decision maker has created an approximating model by a specification search that we do not model. The decision maker believes that data will come from<sup>3</sup> an unknown member of a *set* of unspecified models near his approximating model.<sup>4</sup> 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 *should* 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.<sup>5</sup>

Although the artificial agents within a rational expectations model trust the model, a model's author often doubts it, especially when calibrating it or after performing specification tests. There are several good reasons for wanting to extend rational expectations models to acknowledge fear of model misspecification.<sup>6</sup> First, doing so accepts Muth's (1961) idea of putting econometricians and the agents being modeled on the same footing: because econometricians face specification doubts, the agents inside the model might too.<sup>7</sup> Second, in various contexts, rational expectations models underpredict prices

---

<sup>3</sup> Or, in the case of the robust filtering problems posed in chapter 17, *have* come from.

<sup>4</sup> We say "unspecified" because of how these models are formed as statistical perturbations to the decision maker's approximating model.

<sup>5</sup> In a rational expectations model, each agent's model (i.e., his subjective joint probability distribution over exogenous and endogenous variables) is determined by the equilibrium. It is 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 eliminates all free parameters associated with people's beliefs. For example, see Hansen and Sargent (1980) and Sargent (1981).

<sup>6</sup> In chapter 16, we explore several mappings, the fixed points of which restrict a robust decision maker's approximating model. As is usually the case with rational expectations models, we are silent about the process by which an agent arrives at an approximating model. A qualification to the claim that rational expectations models do not describe the process by which agents form their models comes from the literature on adaptive learning. There, agents who use recursive least squares learning schemes eventually come to know enough to behave as they should in a self-confirming equilibrium. Early examples of such work are Bray (1982), Marcet and Sargent (1989), and Woodford (1990). See Evans and Honkapohja (2001) for new results.

<sup>7</sup> This argument might offend someone with a preference against justifying modeling assumptions on behavioral grounds.

of risk from 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 mispricing.<sup>8</sup> Agents' caution in responding to concerns about model misspecification can raise prices assigned to macroeconomic risks and lead to reinterpreting them as compensation for bearing model uncertainty instead of risks with known probability distributions. 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.<sup>9</sup>

## 1.4. Our extensions of robust control theory

Among ways we adapt and extend robust control theory so that it can be applied to economic problems, six important ones are discounting; a reinterpretation of the “worst-case shock process”; extensions to several multi-agent settings; stochastic interpretations of perturbations to models; a way of calibrating plausible fears of model misspecification as measured by the parameter  $\theta$  in figure 1.1.1; and formulations of robust estimation and filtering problems.

### 1.4.1. Discounting

Most presentations of robustness in control theory treat undiscounted problems, and the few formulations of discounting that do appear differ from the way economists would set things up.<sup>10</sup> In this book, we formulate discounted problems that preserve the recursive structure of decision problems that macroeconomists and other applied economists use so widely.

---

<sup>8</sup> 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 (2004). Barillas, Hansen, and Sargent (2007) argue that Tallarini's and Alvarez and Jermann's measures of the costs of reducing aggregate fluctuations are flawed if what they measure as a market price of risk is instead interpreted as a market price of model uncertainty.

<sup>9</sup> We suspect that his doubts about having a properly specified macroeconomic model explains why, when he formulated comprehensive proposals for the conduct of monetary and fiscal policy, Friedman (1953, 1959) did not use a formal Bayesian expected utility framework, like the one he had used in Friedman and Savage (1948).

<sup>10</sup> Compare the formulations in Whittle (1990) and Hansen and Sargent (1995).



### 1.4.2. Representation of worst-case shock

As we shall see, in existing formulations of robust control theory, shocks that represent misspecification are 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 economists<sup>11</sup> 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 7, we use the “Big  $K$ , little  $k$  trick” from the literature on recursive competitive equilibria to reformulate misspecification perturbations to an approximating model as exogenous processes that cannot be influenced by the decision maker. As we illustrate in the analysis of the permanent income model of chapter 10, this reinterpretation of the worst-case shock process is useful in a variety of economic models.

### 1.4.3. Multiple agent settings

In formulations from the control theory literature, the decision maker’s model of the state transition dynamics is a primitive part of (i.e., an exogenous input into) the statement of the problem. In multi-agent dynamic economic problems, it is not. Instead, 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 equilibrium outcomes. In this book, we describe ways of formulating the decision maker’s approximating model when he and possibly other decision makers are concerned about model misspecification, perhaps to differing extents. We impose a common approximating model on all decision makers, but allow them to express different degrees of mistrust of that model and to have different objectives. As we explain in chapters 12, 15, and 16, this is a methodologically conservative approach that adapts the concept of a Nash equilibrium to incorporate concerns about robustness. The hypothesis of a common approximating model preserves much of the discipline of rational expectations, while the hypothesis that agents have different interests and different concerns about robustness implies a precise sense in which *ex post* they behave as if they had different models. We thereby attain a disciplined way of modeling apparent heterogeneity of beliefs.<sup>12</sup>

---

<sup>11</sup> For example, Christopher Sims expressed this view to us.

<sup>12</sup> Brock and deFontnouvelle (2000) describe a related approach to modeling heterogeneity of beliefs.

#### 1.4.4. *Explicitly stochastic interpretations*

Much of this book is about linear-quadratic problems for which a convenient certainty equivalence result described in chapter 2 permits easy transitions between nonstochastic and stochastic versions of a problem. Chapter 3 describes the relationship between stochastic and nonstochastic setups.

#### 1.4.5. *Calibrating fear of misspecification*

Rational expectations models presume that decision makers know the correct model, a probability distribution over sequences of outcomes. One way to justify this assumption is to appeal to adaptive theories of learning that endow agents with very long histories of data and allow a Law of Large Numbers to do its work.<sup>13</sup> But after observing a short time series, a statistical learning process will typically leave agents undecided among members of a set of models, perhaps indexed by parameters that the data have not yet pinned down well. This observation is the starting point for the way that we use detection error probabilities to discipline the amount of model uncertainty that a decision maker fears after having studied a data set of length  $T$ .

#### 1.4.6. *Robust filtering and estimation*

Chapter 17 describes a formulation of some robust filtering problems that closely resemble problems in the robust control literature. This formulation is interesting in its own right, both economically and mathematically. For one thing, it has the useful property of being the dual of a robust control problem. However, as we discuss in detail in chapter 17, this problem builds in a peculiar form of commitment to model distortions that had been chosen earlier but that one may not want to consider when making current decisions. For that reason, in chapter 18, we describe a class of robust filtering and estimation problems without commitment to those prior distortions. Here the decision maker carries along the density of the hidden states given the past signal history computed under the approximating model, then considers hypothetical changes in this density and in the state and signal dynamics looking forward.

---

<sup>13</sup> For example, see work summarized by Fudenberg and Levine (1998), Evans and Honkapohja (2001), and Sargent (1999a). The justification is incomplete because economies where agents use adaptive learning schemes typically converge to self-confirming equilibria, not necessarily to full rational expectations equilibria. They may fail to converge to rational expectations equilibria because histories can contain an insufficient number of observations about off-equilibrium-path events for a Law of Large Numbers to be capable of eradicating erroneous beliefs. See Cho and Sargent (2007) for a brief introduction to self-confirming equilibria and Sargent (1999a) for a macroeconomic application.

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

Ordinary optimal control theory assumes that decision makers know 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 stem from this association.<sup>14</sup> For example, the Lucas critique (1976) is an application of 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.<sup>15</sup>

Robust control theory alters the mapping from shock temporal properties to decision rules by treating the decision maker's model as an approximation and seeking a *single* rule to use for a *set* of vaguely specified alternative models expressed in terms of distortions to the shock processes in the approximating model. Because they are allowed to feed back arbitrarily on the history of the states, such distortions can represent misspecified 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 decision rules to be functions of parameters that characterize the stochastic processes impinging on agents' constraints. A concern for model misspecification alters these cross-equation restrictions by inspiring the robust decision maker to act as if he had beliefs that seem to *twist* or *slant* probabilities in ways designed to make his decision rule less fragile to misspecification. Formulas presented in chapters 2 and 7 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 appropriately *slanted* laws of motion in the Hansen-Sargent (1980) forecasting formulas. This finding 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

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

---

<sup>14</sup> Stokey and Lucas with Prescott (1989) is a standard reference on using control theory to construct dynamic models in macroeconomics.

<sup>15</sup> See Hansen and Sargent (1980, 1981, 1991).


Where  $y^*$  denotes next period's state vector, let the data truly come from a Markov process with one step transition density  $f(y^*|y)$  that we assume has invariant distribution  $\mu(y)$ . 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  $\alpha$ . 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  $\alpha$  by maximum likelihood. Under some regularity conditions, the maximum likelihood estimator  $\hat{\alpha}_o$  converges in large samples to<sup>16</sup>

$$\text{plim } \hat{\alpha}_o = \operatorname{argmin}_{\alpha \in A} \int I(f_\alpha, f)(y) d\mu(y) \quad (1.6.1)$$

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

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

It can be shown that  $I(f_\alpha, f)(y) \geq 0$ . Figure 1.6.1 depicts how the probability limit  $\hat{\alpha}_o$  of the estimator of the parameters of a misspecified model makes  $I(f_\alpha, f) = \int I(f_\alpha, f)(y) d\mu(y)$  as small as possible. When the model is misspecified, the minimized value of  $I(f_\alpha, f)$  is positive.



**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  $\alpha$  eventually select the misspecified model  $f_{\alpha_o}$  that is closest to  $f$  as measured by entropy  $I(f_\alpha, f)$ .

Sims (1993) and Hansen and Sargent (1993) have used this framework to deduce the consequences of various types of misspecification for estimates of

<sup>16</sup> Versions of this result occur in White (1982, 1994), Vuong (1989), Sims (1993), Hansen and Sargent (1993), and Gelman, Carlin, Stern, and Rubin (1995).

parameters of dynamic stochastic models.<sup>17</sup> For example, they studied the consequences of using seasonally adjusted data to estimate models populated by decision makers who actually base their decisions on seasonally unadjusted data.

## 1.7. Acknowledging misspecification

To study decision making in the presence of model misspecification, we turn the analysis of section 1.6 on its head by taking  $f_{\alpha_o}$  as a given 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. Because he doesn't know  $f$ , a decision maker bases his decisions on the only explicitly specified model available, namely, the misspecified  $f_{\alpha_o}$ . We are silent about the process through which the decision maker discovered his approximating model  $f_{\alpha_o}(y^*|y)$ .<sup>18</sup> We also take for granted the decision maker's parameter estimates  $\alpha_o$ .<sup>19</sup> We impute some doubts about his model to the decision maker. In particular, the decision maker suspects that the data are actually generated by another model  $f(y^*|y)$  with relative entropy  $I(f_{\alpha_o}, f)(y)$ . The decision maker thinks that his model is a good approximation in the sense that  $I(f_{\alpha_o}, f)(y)$  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  $\beta$  and construct the following intertemporal measure of model misspecification:<sup>20</sup>

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

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  $\mathcal{I}(f_{\alpha_o}, f) \leq \eta_0$ , where  $\eta_0$  measures the set of models  $F$  surrounding his approximating model  $f_{\alpha_o}$ . Figure 1.7.1 portrays the decision maker's view of the world. The decision maker wants a single decision rule that is reliable for *all* models  $f$  in the set displayed in figure 1.7.1.<sup>21</sup> This book describes how he

---

<sup>17</sup> Also see Vuong (1989).


<sup>18</sup> See Kreps (1988, chapter 11) for an interesting discussion of the problem of model discovery.

<sup>19</sup> In chapter 9, 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 he wants to be robust.

<sup>20</sup> Hansen and Sargent (2005b, 2007a) provide an extensive discussion of reasons for adopting this measure of model misspecification.

<sup>21</sup> 'Reliable' means good enough, but not necessary optimal, for each member of a set of

can form such a robust decision rule by solving a Bellman equation that tells him how to maximize his intertemporal objective over decision rules when a hypothetical malevolent nature minimizes that same objective by choosing a model  $f$ .<sup>22</sup> That is, we use a max-min decision rule. Positing a malevolent nature is just a device that the decision maker uses to perform a systematic analysis of the fragility of alternative decision rules and to construct a lower bound on the performance that can be attained by using them. A decision maker who is concerned about robustness naturally seeks to construct bounds on the performance of potential decision rules, and the malevolent agent helps the decision maker do that.



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

## 1.8. Why entropy?

To assess the robustness of a decision rule to misspecification of an approximating model requires a way to measure just how good an approximation that model is. In this book, we use the relative entropy to measure discrepancies between models. Of course, relative entropy is not the only way we

---

models. The Lucas critique, or dynamic programming, tells us that it is impossible to find a single decision rule that is *optimal* for all  $f$  in this set. Note how the one-to-one mapping from transition laws  $f$  to decision rules that is emphasized in the Lucas critique depends on the decision maker knowing the model  $f$ . We shall provide a Bayesian interpretation of a robust decision rule by noting that, *ex post*, the max-min decision rule is optimal for *some* model within the set of models.

<sup>22</sup> See Milnor (1951, 1954) for an early formal use of the fiction of a malevolent agent.

could measure discrepancies between alternative probability distributions.<sup>23</sup> But in using relative entropy, we follow a substantial body of work in applied mathematics that reaps benefits from entropy in terms of tractability and interpretability. In particular, using entropy to measure model discrepancies enables us to appeal to the following outcomes:

1. In the general nonlinear case, using entropy to measure model discrepancies means that concerns about model misspecification can be represented in terms of a continuation value function that emerges as the indirect utility function after minimizing the decision maker's continuation value with respect to the transition density, subject to a penalty on the size of conditional entropy. That indirect utility function implies a tractable "risk-sensitivity" adjustment to continuation values in Bellman equations. In particular, we can represent a concern about robustness by replacing  $E_t V(x_{t+1})$  in a Bellman equation with  $-\theta \log E_t \left( \exp \left( \frac{-V(x_{t+1})}{\theta} \right) \right)$ , where  $\theta > \underline{\theta} > 0$  is a parameter that measures the decision maker's concern about robustness to misspecification. (We shall relate the lower bound  $\underline{\theta}$  to  $H_\infty$  control theory in chapter 8.) The simple  $\log E_t \exp$  form of this adjustment follows from the decision to measure model discrepancy in terms of entropy.
2. In problems with quadratic objective functions and linear transition laws, using relative entropy to measure model misspecification leads to a simple adjustment to the ordinary linear-quadratic dynamic programming problem. Suppose that the transition law for the state vector in the approximating model is  $x_{t+1} = Ax_t + Bu_t + C\epsilon_{t+1}$ , where  $\epsilon_{t+1}$  is an i.i.d. Gaussian vector process with mean 0 and identity covariance. Using relative entropy to measure discrepancies in transition laws implies a worst-case model that perturbs the distribution of  $\epsilon_{t+1}$  by enhancing its covariance matrix and appending a mean vector  $w_{t+1}$  that depends on date  $t$  information. Value functions remain quadratic and the distribution associated with the perturbed model remains normal. Because a form of certainty equivalence prevails,<sup>24</sup> it is sufficient to keep track of the mean distortion when solving the control problem. This mean distortion contributes  $.5w_{t+1}' \cdot w_{t+1}$  to the relative entropy discrepancy between the approximating model and the alternative model. As a consequence, a term  $\theta w_{t+1}' w_{t+1}$  is appended to the one-period return function when computing the robust control and a worst-case conditional mean.

---

<sup>23</sup> Bergemann and Schlag (2005) use Prohorov distance rather than entropy to define the set of probability models against which decision makers seek robustness.

<sup>24</sup> See page 33.

3. As we shall see in chapter 9, entropy connects to a statistical theory for discriminating one model from another. The theory of large deviations mentioned in chapter 3 links statistical discrimination to a risk-sensitivity adjustment.<sup>25</sup>

## 1.9. Why max-min?

We answer this question by posing three other questions.

1. What does it mean for a decision rule to be robust? A robust decision rule performs well under the variety of probability models depicted in figure 1.7.1. How might one go about investigating the implications of alternative models for payoffs under a given decision rule? A good way to do this is to compute a lower bound on value functions by assessing the *worst* performance of a given decision rule over a range of alternative models. This makes max-min a useful tool for searching for a robust decision rule.
2. Instead of max-min, why not simply ask the decision maker to put a prior distribution over the set of alternative models depicted in figure 1.7.1? Such a prior would, in effect, have us form a new model – a so-called hypermodel – and thereby eliminate concerns about the misspecification of *that* model. Forming a hypermodel would allow the decision maker to proceed with business as usual, albeit with what may be a more complex model and a computationally more demanding control problem. We agree that this “model averaging” approach is a good way to address some well-structured forms of model uncertainty. Indeed, in chapter 18 we shall use model averaging and Bayesian updating when we study problems that call for combined estimation and control. But the set of alternative models can be so vast that it is beyond the capacity of a decision maker to conjure up a unique well behaved prior. And even when he can, a decision maker might also want decisions to be robust to *whatever* prior he could imagine over this set of models.

More is at issue than the choice of the prior distribution to assign to distinct well specified models. The specification errors that we fear might be more complex than can be represented with a simple model averaging approach. It is reasonable to take the view that each of the distinct models being averaged is itself an approximation. The decision maker might lack precise ideas about how to describe the alternative specifications that worry him and about how to form prior distributions over

---

<sup>25</sup> Anderson, Hansen, and Sargent (2003) extensively exploit these connections.



them. Perhaps he can't articulate the misspecifications that he fears, or perhaps the set of alternative models is too big to comprehend.<sup>26</sup>

Our answer to this second question naturally leads to a reconsideration of the standard justification for being a Bayesian.

3. "Why be a Bayesian?" Savage (1954) gave an authoritative answer by describing axioms that imply that a rational person can express all of his uncertainty in terms of a unique prior. However, Schmeidler (1989) and Gilboa and Schmeidler (1989) altered one of Savage's axioms to produce a model of what it means to be a rational decision maker that differs from Savage's Bayesian model. Gilboa and Schmeidler's rational decision maker has multiple priors and behaves as a max-min expected utility decision maker: the decision maker *maximizes* and assumes that nature chooses a probability to *minimize* his expected utility. We are free to appeal to Gilboa and Schmeidler's axioms to rationalize the form of max-min expected utility decision making embedded in the robust control theories that we study in this book.<sup>27</sup>

### 1.10. Is max-min too cautious?

*Our doubts are traitors, And make us lose the good we oft might win, By fearing to attempt.*

— William Shakespeare, *Measure for Measure*, act 1 scene 4

Our use of the detection error probabilities of chapter 9 to restrict the penalty parameter  $\theta$  in figure 1.1.1 protects us against the objection that the max-min expected utility theory embedded in robust control theory is too cautious because, by acting as if he believed the worst-case model, the decision maker puts too much weight on a "very unlikely" scenario.<sup>28</sup> We choose  $\theta$  so that the entropy ball that surrounds the decision maker's approximating model in

---

<sup>26</sup> See Sims (1971) and Diaconis and Freedman (1986) for arguments that forming an appropriate prior is difficult when the space of submodels and the dimensions of parameter spaces are very large.

<sup>27</sup> Hansen and Sargent (2001) and Hansen, Sargent, Turmuhambetova, and Williams (2006) describe how stochastic formulations of robust control "constraint problems" can be viewed in terms of Gilboa and Schmeidler's max-min expected utility model. Interesting theoretical work on model ambiguity not explicitly connected to robust control theory includes Dow and Werlang (1994), Ghirardato and Marinacci (2002), Ghirardato, Maccheroni, and Marinacci (2004), Ghirardato, Maccheroni, Marinacci, and Siniscalchi (2003), and Rigotti and Shannon (2003, 2005), and Strzalecki (2007).

<sup>28</sup> Bewley (1986, 1987, 1988), Dubra, Maccheroni, and Ok (2004), Rigotti and Shannon (2005), and Lopomo, Rigotti, and Shannon (2004) use an alternative to the max-min expected utility model but still one in which the decision maker experiences ambiguity about models. In their settings, incomplete preferences are expressed in terms of model ambiguity

figure 1.7.1 has the property that the perturbed models on and inside the ball are difficult to distinguish statistically from the approximating model with the amount of data at hand. This way of calibrating  $\theta$  makes the likelihood function for the decision maker's worst-case model fit the available data almost as well as his approximating model. Moreover, by inspecting the implied worst-case model, we can evaluate whether the decision maker is focusing on scenarios that appear to be too extreme.

### 1.11. Aren't you just picking a plausible prior?

By interchanging the order in which we maximize and minimize, chapter 7 describes an *ex post* Bayesian interpretation of a robust decision rule.<sup>29</sup> Friendly critics have responded to this finding by recommending that we view robust control as simply a way to select a plausible prior in an otherwise standard Bayesian analysis.<sup>30</sup> Furthermore, one can regard our chapter 9 detection error probability calculations as a way to guarantee that the prior is plausible in light of the historical data record at the disposal of the decision maker.

We have no objection to this argument in principle, but warn the reader that issues closely related to the Lucas (1976) critique mean that it has to be handled with care, as in any subjectivist approach. Imagine a policy intervention that alters a component of a decision maker's approximating model for, e.g., a tax rate, while leaving other components unaltered. In general, *all* equations of the decision maker's worst-case transition law that emerge from the max-min decision process will vary with such interventions. The dependence of other parts of the decision maker's worst-case model on subcomponents of the transition law for the approximating model that embody the policy experiment reflects the context-specific nature of the decision maker's worst-case model. Therefore, parts of the *ex post* worst-case "prior" that describe the evolution of variables *not* directly affected by the policy experiment will depend on the policy experiment. The sense in which robust control is just a way to pick a plausible prior is subtle.

Another challenge related to the Lucas critique pertains when we apply robust control without availing ourselves of the *ex post* Bayesian interpreta-

---

and there is a status quo allocation that plays a special role in shaping how the decision maker ranks outcomes. Some advocates of this incomplete preferences approach say that they like it partly because it avoids what they say is an undue pessimism that characterizes the max-min expected utility model. See Fudenberg and Levine (1995) for how max-min can be used to attain an interesting convergence result for adaptive learning.

<sup>29</sup> We introduce this argument because it provides a sense in which our robust decision rules are admissible in the statistical decision theoretic sense of being undominated.

<sup>30</sup> Christopher A. Sims has made this argument on several occasions.

tion. Throughout this book, whenever we consider changes in the economic environment, we imitate rational expectations policy analysis by imputing common approximating models, one before the policy change, the other after, to all agents in the model and the econometrician (e.g., see chapter 14). It is natural to doubt whether decision makers would fully trust their statistical models after such policy changes.

## 1.12. Why not learn the correct specification?

For much of this book, but not all, we attribute an enduring fear of misspecification to our decision maker. Wouldn't it be more realistic to assume that the decision maker learns to detect and discard bad specifications as data accrue?

One good answer to this question is related to some of the points made in section 1.9. In chapter 9, we suggest calibrating the free parameter  $\theta$  borne by the "gentleman" in the bottom panel of figure 1.1.1 so that, even with nondogmatic priors, it would take long time series to distinguish among the alternative specifications about which the decision maker is concerned. Because our decision maker discounts the future, he cannot avoid facing up to his model specification doubts simply by waiting for enough data.<sup>31</sup> Thus, one answer is that, relative to his discount factor, it would take a long time for him to learn not to fear model misspecification.

However, we agree that it is wise to think hard about what types of misspecification fears you can expect learning to dispel in a timely way, and which types you cannot. But what are good ways to learn when you distrust your model? Chapters 17 and 18 are devoted to these issues.<sup>32</sup> We present alternative formulations of robust estimation and filtering problems and suggest ways to learn in the context of distrusted approximating models. Our approach allows us to distinguish types of model misspecification fears that a decision maker can eventually escape by learning from types that he cannot.<sup>33</sup>

---

<sup>31</sup> As we shall see, one reason that it takes a very long data set to discriminate between the models that concern the decision maker is that often they closely approximate each other at high frequencies and differ mostly at very low frequencies. Chapter 8 studies robustness from the viewpoint of the frequency domain.

<sup>32</sup> Also see Hansen and Sargent (2005b, 2007a, 2007b).

<sup>33</sup> Epstein and Schneider (2006) also make this distinction. In the empirical model of Hansen and Sargent (2007b), a representative consumer's learning within the sample period reduces his doubts about the distribution of some unknown parameters, but does little to diminish his doubts about the distribution over difficult to distinguish submodels, one of which confronts him with long-run risk in the growth rate of consumption.

### 1.13. Is the set of perturbed models too limited?

Parts of this book are devoted to analyzing situations in which the decision maker's approximating model and the statistical perturbations to it that bother him all take the form of the stochastic linear evolution

$$x_{t+1} = Ax_t + Bu_t + C(\epsilon_{t+1} + w_{t+1}) \quad (1.13.1)$$

where  $x_t$  is a state vector,  $u_t$  a control vector,  $\epsilon_{t+1}$  an i.i.d. Gaussian shock with mean 0 and covariance  $I$ , and  $w_{t+1}$  is a vector of perturbations to the mean of  $\epsilon_{t+1}$ . Under the approximating model,  $w_{t+1} = 0$ , whereas under perturbed models,  $w_{t+1}$  is allowed to be nonzero and to feed back on the history of past  $x_t$ 's.

Some critics have voiced the complaint that this class of perturbations excludes types of misspecified dynamics that ought to concern a decision maker, such as unknown parameter values, misspecification of higher moments of the  $\epsilon_{t+1}$  distribution, and various kinds of "structured uncertainty." We think that this complaint is misplaced for the following reasons:

1. For the problems with quadratic objective functions and approximating models like (1.13.1) with  $w_{t+1} = 0$ , restricting ourselves to perturbations of the form (1.13.1) turns out not to be as restrictive as it might at first seem. In chapters 3 and 7, we permit a much wider class of alternative models that we formulate as absolutely continuous perturbations to the transition density of state variables. We show that when the decision maker's objective function is quadratic and his approximating model is linear with Gaussian  $\epsilon_{t+1}$ , then he chooses a worst-case model that is of the form (1.13.1) with a  $C$  that is usually only slightly larger and a  $w_{t+1}$  that is a linear function of  $x_t$ . We shall explain why he makes little or no error by ignoring possible misspecification of the volatility matrix  $C$ .
2. In section 19.2 of chapter 19, we show how more structured kinds of uncertainty can be accommodated by slightly reinterpreting the decision maker's objective function.
3. When the approximating model is a linear state evolution equation with Gaussian disturbances and the objective function is quadratic, worst case distributions are also jointly Gaussian. However, making the approximating model be non-Gaussian and non-linear or making the objective function be not quadratic leads to non-Gaussian worst-case joint probability distributions, as chapter 3 indicates. Fortunately, by extending the methods of chapters 17 and 18, as Hansen and Sargent (2005, 2007a) do, we know how to model robust decision makers who learn about non-linear

models with non-Gaussian shock distributions while making decisions. The biggest hurdles in carrying out quantitative analyses like these are computational. Most of the problems studied in this book are designed to be easy computationally by staying within a linear-quadratic-Gaussian setting. But numerical methods allow us to tackle analogous problems outside the LQG setting.<sup>34</sup>

### 1.14. Is robust control theory positive or normative?

Robust control and estimation theory has both normative and positive economic applications. In some contexts, we take our answer to question (2) in the preceding section to justify a positive statement about how people *actually* behave. For example, we use this interpretation when we apply robust control and estimation theory to study asset pricing puzzles by constructing a robust representative consumer whose marginal evaluations determine market prices of risk (see Hansen, Sargent, and Tallarini (1999), Hansen, Sargent, and Wang (2002), and chapter 13).

Monetary policy authorities and other decision makers find themselves in situations where their desire to be cautious with respect to fears of model misspecification would inspire them to use robust control and estimation techniques.<sup>35</sup> Normative uses of robust control theory occur often in engineering.

### 1.15. Other lessons

Our research program of refining typical rational expectations models to attribute specification doubts to the agents inside of them has broadened our own understanding of rational expectations models themselves. Struggling with the ideas in this book has taught us much about the structure of recursive models of economic equilibria,<sup>36</sup> the relationship between control and estimation problems, and Bayesian interpretations of decision rules in dynamic rational expectations models. We shall use the macroeconomist's Big  $K$ , little  $k$  trick with a vengeance.

The 1950s-1960s control and estimation theories lampooned in the top panel of figure 1.1.1 have contributed enormously to the task of constructing dynamic equilibrium models in macroeconomics and other areas of applied economic dynamics. We expect that the robust control theories represented

---

<sup>34</sup> See Cogley, Colacito, Hansen, and Sargent (2007) for an example.

<sup>35</sup> Blinder (1998) expresses doubts about model misspecification that he had when he was vice chairman of the Federal Reserve System and how he coped with them.

<sup>36</sup> For example, see chapter 12.

in the bottom panel of that figure will also bring many benefits that we cannot anticipate.

### **1.16. Topics and 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, unlike the control literature, which focuses on undiscounted problems, we formulate discounted problems. Incorporating discounting involves substantial work, especially in chapter 8, 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 12 and 13); (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 the leader fears model misspecification (chapter 16). 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, in which a pervasive certainty equivalence principle allows a nonstochastic presentation of most of the control and filtering theory.

This book is organized as follows. Chapter 2 summarizes a set of practical results at a relatively nontechnical level. A message of this chapter is that although sophisticated arguments from chapters 7 and 8 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 chapter 3. Chapter 3 tells how the key ideas about robustness generalize to models that are not linear quadratic.

Chapters 4 and 5 are about optimal control and filtering when the decision maker trusts his model. These chapters contain a variety of useful results for characterizing the linear dynamic systems that are widely used in macroeconomics. Chapter 4 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 equilibrium models. Later chapters apply these methods to various problems: to compute robust decision rules as solutions of two-player zero-sum games; to compute robust filters via another two-player zero-sum game; and to compute equilibria of robust Stackelberg or Ramsey problems in macroeconomics. Chapter 5 emphasizes that the Kalman filter is the *dual* (in a sense familiar to economists from their use of Lagrange multipliers) of the basic linear-quadratic dynamic programming problem of chapter 4 and sets the stage for a related duality result for a robust filtering problem to be presented in chapter 17.

The remaining chapters are about making wise decisions when a decision maker distrusts his model. Within a one-period setting, chapter 6 introduces two-player 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 6 is a good one for addressing some important conceptual issues. In particular, in this chapter we state multiplier and constraint problems, two different two-player zero-sum games that induce robust decision rules. We use the Lagrange multiplier theorem to connect the problems.

Chapters 7 and 8 extend and modify results in the control literature to formulate robust control problems with discounted quadratic objective functions and linear transition laws. Chapter 7 represents things in the time domain, while chapter 8 works in the frequency domain. Incorporating discounting requires carefully restating the control problems used to induce robust decision rules. Chapters 7 and 8 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 two-player zero-sum games in (1). In category (1), we present a detailed account of several two-player 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 express concerns about model misspecification. Two of them are expressed in the frequency domain: the  $H_\infty$  and entropy criteria. The entropy objective function summarizes model specification doubts with a single parameter. That parameter relates to a Lagrange multiplier in a two-player zero-sum constraint game, and

also to the *risk-sensitivity* parameter of Jacobson (1973) and Whittle (1990), as modified for discounting by Hansen and Sargent (1995).

Chapters 7 and 8 show how robustness is induced by using max-min strategies: the decision maker maximizes while nature minimizes over a set of models that are close to the approximating model. There are alternative timing protocols in terms of which a two-player zero-sum game can be cast. A main finding of chapter 7 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 also facilitates computing equilibria in multiple-agent economics.

Arthur Goldberger and Robert E. Lucas, Jr., warned applied economists to beware of theorists bearing free parameters (see figure 1.1.1). Relative to settings in which decision makers completely trust their models, the multiplier and constraint problems of chapters 7 and 8 each bring one new free parameter that expresses a concern about model misspecification,  $\theta$  for the multiplier problem and  $\eta$  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 9 proposes a way to calibrate these parameters by using the statistical theory for discriminating models.<sup>37</sup> We apply this theory in chapters 10 and 14.

Chapter 10 uses the permanent income model of consumption as a laboratory for illustrating some of the concepts from chapters 7 and 8. 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 10 as a laboratory for asset pricing in chapter 13. But first, chapters 11 and 12 describe how to decentralize the solution of a planning problem with a competitive equilibrium. Chapter 11 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, another with sequential trading of one-period Arrow securities. In that sequential setting, we give a recursive representation of equilibrium prices. Chapter 11 describes a setting where the representative agent has no concern about model misspecification, while chapter 12 extends the characterizations of chapter 11 to situations where the representative decision maker fears model misspecification.

Chapter 13 builds on the chapter 12 results to show how fear of model

---

<sup>37</sup> See Anderson, Hansen, and Sargent (2003).



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. The adjustment measures the representative consumer's fear that the approximating model is misspecified. The adjustment for robustness resembles ones that financial economists use 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 10. A remarkable observational equivalence result identifies a locus of pairs of discount factors and robustness multipliers, all of which imply identical real allocations.<sup>38</sup> Nevertheless, prices of risky assets vary substantially across these pairs. In chapter 14, we revisit some quantitative findings of Tallarini (2000) and reinterpret asset pricing patterns that he imputed to very high risk aversion in terms of a plausible fear of model misspecification. We measure a plausible fear of misspecification by using the detection error probabilities introduced in chapter 9.

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. Important aspects of that common model, those governing endogenous state variables, are equilibrium outcomes. 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 this model governs the data.<sup>39</sup> To preserve that empirical power in an equilibrium with multiple decision makers who fear model misspecification, we impose that all decision makers share a common approximating model.<sup>40</sup> The model components that describe endogenous state variables are equilibrium outcomes that depend on agents' robust decision making processes, i.e., on the solutions to their max-min 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

---

<sup>38</sup> This result establishes a precise sense in which, so far as real quantities are concerned, increased fear of model misspecification acts just like *reduced* discounting of the future, so that its effects on real quantities can be offset by *increasing* the rate at which future payoffs are discounted.

<sup>39</sup> The restriction that they share a common model is the feature that makes free parameters governing expectations disappear. This is what legitimizes a law of large numbers that underlies rational expectations econometrics.

<sup>40</sup> 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 modeling, namely, that the decision makers' approximating model actually *does* generate the data.

approximating model and 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 solving pairs of 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 differ because their objectives differ. In this sense, the model produces endogenous *ex post* heterogeneity of beliefs.

In chapter 16, we alter the timing protocol to study a 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 4.

Chapter 17 extends the analysis of filtering from chapter 5 by describing a robust filtering problem that is dual to the control problem of chapter 7.<sup>41</sup> This recursive filtering problem requires that a time  $t$  decision maker must respect distortions to the distribution of the hidden state that he inherits from past decision makers. As a consequence, in this problem, bygones are not bygones:<sup>42</sup> the decision maker's concerns about *past* returns affect his estimate of the current value of a hidden state vector.

Chapter 18 uses a different criterion than chapter 17 and finds a different robust filter. We think that the chapter 18 filter is the appropriate one for many problems and give some examples. The different filters that emerge from chapters 17 and 18 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 this book.

Chapter 19 concludes by confronting some of the confining aspects of our work, some criticisms that we have heard, and opportunities for further progress.

---

<sup>41</sup> We originally found this problem by stating and solving a conjugate problem of a kind familiar to economists through duality theory. By faithfully following where duality leads, we discovered a filtering problem that is peculiar (but not necessarily uninteresting) from an economic standpoint. A sketch of this argument is presented in appendix A of chapter 17.

<sup>42</sup> But see the epigraph from William Stanley Jevons quoted at the start of chapter 18.

## 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 Millennium*, 1996

#### 2.1. Introduction

A model maps a sequence of decisions into a sequence of outcomes. 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 a 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.<sup>1</sup> After describing possible misspecifications as a set of perturbations to an approximating model, 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 contexts that we exploit extensively in subsequent chapters. This chapter concentrates mainly on single-agent decision theory, but chapters 11, 15, and 16 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 extend the notion of a rational expectations equilibrium to situations in which decision makers have different amounts of confidence in a common approximating model.<sup>2</sup> Chapter 3

---

<sup>1</sup> Later chapters provide technical details that justify assertions made in this chapter.

<sup>2</sup> Chapter 11 discusses competitive equilibria in representative agent economies; chapter 15 injects motives for robustness into Markov perfect equilibria for two-player dynamic games; and chapter 16 studies Stackelberg and Ramsey problems. In Ramsey problems, a 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.

studies models with more general return and transition functions and shows that many of the insights of this chapter apply beyond the linear-quadratic setting. The LQ setting is computationally tractable, but also reveals most of the conceptual issues that apply with more general functional forms.

## 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

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

where  $\{\check{\epsilon}_t\}$  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  $\check{\epsilon}_{t+1}$  can represent only a very limited class of approximation errors and in particular cannot depict such examples of misspecified *dynamics* as are represented in models with nonlinear and time-dependent feedback of  $y_{t+1}$  on past states. To represent dynamic misspecification,<sup>3</sup> we surround (2.2.1) with a set of alternative models of the form

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

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  $y$

$$w_{t+1} = g_t(y_t, y_{t-1}, \dots), \quad (2.2.3)$$

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) were conditionally distributed as  $\mathcal{N}(w_{t+1}, I)$  rather than as  $\mathcal{N}(0, I)$ . Thus, we capture the idea

---

<sup>3</sup> In chapters 3 and 6, we allow a broader class of misspecifications. Chapter 3 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 shock in the distorted model that turns out not to affect the distortion to the conditional mean of the shock. In many applications, the adjustment to the covariance matrix is quantitatively insignificant. It vanishes in the case of continuous time. See Anderson, Hansen, and Sargent (2003) and Hansen, Sargent, Turmuhambetova, and Williams (2006).

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

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

where  $E_t$  denotes mathematical expectation evaluated with model (2.2.2) and conditioned on  $y^t = [y_t, \dots, y_0]$ . In section 2.3, chapter 3, and chapter 9, we shall interpret the left side of (2.2.4) as a statistical measure of the discrepancy between the distorted and approximating models.

The alternative models differ from the approximating model by having shock processes whose conditional means are not zero and that can feed back in potentially complicated ways on the history of the state. Notice that our specification leaves the conditional volatility of the shock, as parameterized by  $C$ , unchanged. We adopt this specification for computational convenience. We show in chapter 3 the useful result that our calculations for a worst-case conditional mean  $w_{t+1}$  remain unaltered when we also allow conditional volatilities  $C$  to differ in the approximating and perturbed 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).<sup>4</sup> The decision maker forsakes learning to improve his specification because  $\eta_0$  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 9.<sup>5</sup>

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 misspecification of the approximating model.

We compute robust decision 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 (also known as 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 7 and 8, equilibrium outcomes and decision rules for the games

---

<sup>4</sup> See chapter 3 for a specification of the approximating model as a joint probability density over an infinite sequence of  $y_t$ s and misspecifications that are represented as alternative joint probability densities.

<sup>5</sup> However, see chapter 18 and Hansen and Sargent (2005b, 2007a) for ways to include robust forms of learning.

coincide, a consequence of the zero-sum feature of all of the games.<sup>6</sup> This makes the games easy to solve. Computing robust decision rules comes down to solving Bellman equations for dynamic programming problems that are very similar to equations 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.

### 2.2.1. Dynamic programming without model misspecification

The standard dynamic programming problem assumes that the transition law is correct.<sup>7</sup> Let the one-period loss function be  $r(y, u) = -(y'Qy + u'Ru)$  where the matrices  $Q$  and  $R$  are symmetric and together with  $A$  and  $B$  in (2.2.1) satisfy some stabilizability and detectability assumptions set forth in chapter 4. The *optimal linear regulator problem* is

$$-y_0'Py_0 - p = \max_{\{u_t\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t r(y_t, u_t), \quad 0 < \beta < 1, \quad (2.2.5)$$

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  $\tilde{\epsilon}$ , 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

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

where the maximization is subject to

$$y^* = Ay + Bu + C\tilde{\epsilon}, \quad (2.2.7)$$

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

Subject to assumptions about  $A, B, R, Q, \beta$  to be described in chapter 4, some salient facts about the optimal linear regulator are the following:

---

<sup>6</sup> The zero-sum feature perfectly misaligns the preferences of the two players and thereby renders timing protocols irrelevant. See chapter 7 for details.

<sup>7</sup> Many technical results and computational methods for the linear quadratic problem without concerns about robustness are catalogued in chapter 4.

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

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

The optimal decision rule is  $u_t = -F y_t$  where

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

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  $p = \frac{\beta}{1-\beta} \text{trace} P C C'$ . 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  $C$  by setting  $E \check{\epsilon}_t \check{\epsilon}_t' = I$ . Therefore, the matrix  $C$  determines the covariance matrix  $C C'$  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  $u_t = -F y_t$  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 33.
3. *Shadow prices.* Since the value function is  $-y_0' P y_0 - p$ , the vector of shadow prices of the initial state is  $-2P y_0$ . 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  $(y_t, \mu_t)$ . The optimal policy solves this difference equation subject to an initial condition for  $y_0$  and a transversality or detectability condition  $E_0 \sum_{t=0}^{\infty} \beta^t r(y_t, u_t) > -\infty$ . In chapter 4, we show that subject to these boundary conditions, the vector difference equation consisting of the first-order conditions is solved by setting  $\mu_t = P y_t$ , where  $P$  solves the Riccati equation (2.2.8).

### 2.3. Measuring model misspecification with 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 33, 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  $u_t = -Fy_t$ . Let  $A_o = A - BF$  and write the approximating model as

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

and a distorted model as<sup>8</sup>

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

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

$$\check{\epsilon}_{t+1} = \epsilon_{t+1} + w_{t+1}, \quad (2.3.3)$$

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 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

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

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.<sup>9</sup>

---

<sup>8</sup> Chapter 3 allows a larger set of perturbations to the approximating model and gives an appropriate definition of entropy.

<sup>9</sup> Define the likelihood ratio  $m(f(y^*|y)) = \frac{f(y^*|y)}{f_o(y^*|y)}$ . Then notice that

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

where the subscript  $f_o$  means integration with respect to the approximating model  $f_o$ . Hansen and Sargent (2005b, 2007a) exploit such representations of entropy. See chapter 3.



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 is<sup>10</sup>

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

where both  $u$  and  $w$  are measurable functions of  $y^t$ . In subsection 3.11 of chapter 3, we verify that the expected log-likelihood is

$$I(w_{t+1}) = .5w'_{t+1}w_{t+1}. \quad (2.3.4)$$

In chapter 9, we describe how measures like (2.3.4) govern the distribution of test statistics for discriminating among models. In chapter 13, we show how the log-likelihood ratio 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

$$R(w) = 2E_0 \sum_{t=0}^{\infty} \beta^{t+1} I(w_{t+1}), \quad (2.3.5)$$

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,

$$R(w) \leq \eta_0. \quad (2.3.6)$$

In the next section, we construct decision rules that work well over a set of models that satisfy (2.3.6). 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.

---

<sup>10</sup> In a continuous-time diffusion setting, Hansen, Sargent, Turmuhambetova, and Williams (2006) 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).

## 2.4. Two robust control problems

This section states two robust control problems: a constraint problem and a multiplier problem. The two problems differ in how they treat constraint (2.3.6). Under appropriate 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  $\eta_0$  satisfying  $\bar{\eta} > \eta_0 \geq 0$ , a constraint problem is

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

where the extremization<sup>11</sup> is subject to the distorted model (2.2.2) and the entropy constraint (2.3.6), 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

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

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 (8.4.8) for  $\underline{\theta}$  shows how the value of  $\underline{\theta}$  depends on the return function, the discount factor, and the transition law. Tests for whether  $\theta > \underline{\theta}$  are presented in formula (7.9.1) and in chapter 8, especially section 8.7. We shall discuss the lower bound  $\underline{\theta}$  and an associated upper bound  $\bar{\eta}$  extensively in chapter 8.

---

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

Chapters 7 and 8 state conditions on  $\theta$  and  $\eta_0$  under which the two problems have identical solutions, namely, decision rules  $u_t = -Fy_t$  and  $w_{t+1} = Ky_t$ . Chapter 7 establishes many useful facts about distinct versions of the multiplier problem that employ alternative timing protocols<sup>12</sup> and that justify solving the multiplier problem recursively. Let  $-y_0'Py_0 - p$  be the value of problem (2.4.2). It satisfies the Bellman equation<sup>13</sup>

$$-y'Py - p = \max_u \min_w E \{r(y, u) + \theta\beta w'w - \beta y^{*'}Py^* - \beta p\} \quad (2.4.3)$$

where the extremization is subject to

$$y^* = Ay + Bu + C(\epsilon + w) \quad (2.4.4)$$

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 (2.4.4). Before describing how robust decision rules emerge from the two-player zero-sum game (2.4.2), we mention a kind of certainty equivalence that applies to the multiplier problem.

#### 2.4.1. Modified certainty equivalence principle

On page 29, 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 another certainty equivalence principle 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  $\epsilon_{t+1}$  from the state-transition

---

<sup>12</sup> 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.

<sup>13</sup> In chapter 7, 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  $u_t$ .

equation and  $p$  from the value function  $-y'Py - p$ , without affecting formulas for the decision rules.<sup>14</sup> 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 29 — 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.<sup>15</sup> This effect is featured in the precautionary savings model of chapter 10, a simple version of which we shall sketch in section 2.8.

## 2.5. Robust linear regulator

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

$$y_{t+1} = Ay_t + Bu_t + Cw_{t+1} \quad (2.5.1)$$

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:

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

where the extremization is subject to (2.5.1) and  $y_0$  is given. Let  $-y'_0 P y_0$  be the value of (2.5.2). It satisfies the Bellman equation<sup>16</sup>

$$-y'Py = \max_u \min_w \{r(y, u) + \theta \beta w'w - \beta y^* P y^*\} \quad (2.5.3)$$

---

<sup>14</sup> The certainty equivalence principle stated here shares with the one on page 29 the facts that  $P$  can be computed before  $p$ ; it diverges from the certainty equivalence principle without robustness on page 29 because now  $P$  and therefore  $F$  both depend on the volatility matrix  $C$ . See Hansen and Sargent (2005a) for a longer discussion of certainty equivalence in robust control problems.

<sup>15</sup> 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 3).

<sup>16</sup> Notice how this is a special case of (2.4.3) with  $p = 0$ . The modified certainty equivalence principle implies that the same matrix  $P$  solves (2.5.3) and (2.4.3).

where the extremization is subject to the *distorted* model

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

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<sup>17</sup>

$$-y^{*\prime} \mathcal{D}(P) y^* = -(x' A' + u' B') \mathcal{D}(P) (Ax + Bu) = \min_w \{ \theta w' w - y^{*\prime} P y^* \} \quad (2.5.5)$$

where the minimization is subject to the transition law  $y^* = Ay + Bu + Cw$ . From the minimization problem on the right of (2.5.5), it follows that<sup>18</sup>

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

The Bellman equation (2.5.3) can then be represented as

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

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.<sup>19</sup>

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):

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

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  $u = -Fy$ , where

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

<sup>17</sup> See page 168, item 1, for more details.

<sup>18</sup> Before computing  $\mathcal{D}$  in formula (2.5.5), we always check whether the matrix being inverted on the right side of (2.5.6) is positive definite. This amounts to a check that  $\theta$  exceeds the “breakdown point”  $\underline{\theta}$ .

<sup>19</sup> 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 3.

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

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

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

$$u_t = -Fy_t \quad (2.5.11a)$$

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

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

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

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

$$-\sum_{t=0}^{\infty} \beta^t [y_t' Q y_t + u_t' R u_t]$$

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  $\theta, \beta, Q, R, A, B$ , and  $C$ . Equation (2.5.12) incorporates how the distortion  $w$  feeds 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 worst-case model.* 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 7, we show that the robust decision rule  $u_t = -Fy_t$  solves the ordinary linear regulator problem

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

subject to the distorted transition law

$$y_{t+1} = Ay_t + Bu_t + Cw_{t+1} \quad (2.5.14a)$$

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

$$Y_{t+1} = A^*Y_t \quad (2.5.14c)$$

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  $y_t$ . This is a version of the macroeconomist's Big  $K$ , little  $k$  trick, where  $Y$  plays the role of Big  $K$ . The solution of (2.5.13), (2.5.14) has the outcome that  $Y_t = y_t \forall t \geq 0$ .<sup>20</sup> Chapters 7 and 8 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 8, this formulation helps us 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*, a belief that is distorted relative to the approximating model.<sup>21</sup> This observation is reminiscent of some ideas of Fellner.

3. *Relation to Fellner (1965).* In the introduction to *Probability and Profit*, William Fellner wrote:

---

<sup>20</sup> 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.

<sup>21</sup> A decision rule is said to have a Bayesian interpretation if it is undominated in the sense of being optimal for some model. See Robert (2001, pp. 74-77) and Blackwell and Girschik (1954).

“... the central problems of decision theory may be described as 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 directly observable decision weights (expectation 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 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 of the slanting may even depend on the magnitude of the prize which is at stake when a prospect is being weighted.”

Robust control theory embodies 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  $\beta, R, Q$  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 7 on page 170, we show that equation (2.5.7) implies that

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

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  $y^* = A_F y$ .<sup>22</sup> 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  $\theta$  times the measure of model misspecification  $w'w$ . The parameter  $\theta$  influences the conservative-adjustment operator  $\mathcal{D}$  and also determines the rate at which the bound deteriorates with misspecification. Lowering  $\theta$  lowers the rate at which the bound deteriorates with misspecification. Thus, (2.5.15) provides a sense in which lower values of  $\theta$  provide more conservative estimates of continuation utility and therefore more robust guides to decision making.

---

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



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 7 describes the relationships among these games and the remarkable fact that they have identical outcomes. The analysis of chapter 7 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 by

$$y_{t+1} = Ay_t - BFy_t + Ky_t. \quad (2.5.16)$$

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 way to help 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 occasionally 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

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

We obtain (2.5.17) from (2.5.16) by replacing the worst-case shock  $Ky_t$  with 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.<sup>23</sup> In chapters 13 and 14, 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  $\theta$  and  $H_\infty$  control.* Starting from  $\theta = +\infty$ , lowering  $\theta$  increases the fear of misspecification by lowering the shadow price on the

---

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

norm of the control of the minimizing player. We shall see in chapter 8 that there is a lower bound for  $\theta$ . This lower bound is associated with the *largest* set of alternative models, as measured by entropy, against which it is feasible to seek a robust rule: for values of  $\theta$  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  $\theta$  because it is affiliated with a rule that is robust to the biggest allowable set of misspecifications. We describe the associated  $H_\infty$  control theory in chapter 8. However, the applications that we are interested in usually call for values of  $\theta$  that exceed the cutoff value by far. We explain why in chapter 9, where we use detection error probabilities to discipline the setting for  $\theta$ .

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  $u_t = -Fy_t$  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

$$y_{t+1} = Ay_t + Bu_t + C\epsilon_{t+1} \quad (2.5.18)$$

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

$$U_t = r(y_t, u_t) + \beta \mathcal{R}_t(U_{t+1}) \quad (2.5.19)$$

where

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

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

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

where the maximization is subject to  $y^* = Ay + Bu + C\epsilon$  and  $\epsilon$  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

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

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

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

Evidently, the fixed point  $P$  satisfies  $P = T \circ \mathcal{D}(P)$ , and therefore it 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  $p$  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. For more discussion of these points, see chapter 3.

## 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 it may at first appear. In chapter 3, 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 when return functions are quadratic and the transition law implied by the approximating model is linear. 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  $f$  relative to the approximating model  $f_o$  is

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

Entropy  $I(f_o, f)(y)$  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:

$$-y'Py - p = \max_u \min_f E \{ r(y, u) + 2\theta\beta I(f_o, f)(y) - \beta y^*Py^* - \beta p \}. \quad (2.6.1)$$

Let  $\sigma = -\theta^{-1}$  and consider the inner minimization problem, assuming that  $u = -Fy$ . In chapter 3, we shall show that the extremizing  $f$  is the Gaussian distribution

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

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

$$\hat{C}\hat{C}' = C(I + \sigma C'PC)^{-1}C', \quad (2.6.3)$$

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.<sup>24</sup>

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  $f(y^*|y)$ , 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 13, we use some of these ideas to price assets under alternative assumptions about the set of models against which decision makers seek robustness.

---

<sup>24</sup> 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 (2006) and Anderson, Hansen, and Sargent (2003).

## 2.7. A simple algorithm

Chapter 7 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 7 and 8 describe necessary technical conditions, including restrictions on the magnitude of the multiplier parameter  $\theta$ .<sup>25</sup>

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  $w$  and 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.

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

$$\tilde{B} = [B \quad C] \quad (2.7.1a)$$

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

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

Let  $\text{ext}$  denote extremization – maximization with respect to  $u$ , minimization with respect to  $w$ . The Bellman equation can be written as

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

where the extremization is subject to

$$y^* = Ay + \tilde{B}\tilde{u}. \quad (2.7.3)$$

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

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

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

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

---

<sup>25</sup> The Matlab program `olrprobust.m` described in the appendix implements this algorithm; `doublex9.m` implements a doubling algorithm of the kind described in chapter 4 and Hansen and Sargent (2008); please note that `doublex9.m` solves a *minimum* problem and that  $-\theta^{-1} \equiv \sigma < 0$  connotes a fear of model misspecification.

Partitioning  $\tilde{F}$ , we have

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

$$w_{t+1} = Ky_t. \tag{2.7.6b}$$

The decision rule  $u_t = -Fy_t$  is the robust rule. As mentioned above,  $w_{t+1} = Ky_t$  provides the  $\theta$ -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 4.

### 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).<sup>26</sup> Hansen and Sargent's (1995) discounted version of the risk-sensitive preferences of Jacobson (1973) and Whittle (1990) also uses the  $\mathcal{D}$  operator.

## 2.8. Robustness and discounting in a permanent income model

This section illustrates aspects of robust control theory in the context of a linear-quadratic version of a simple permanent income model.<sup>27</sup> 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. 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.

---

<sup>26</sup> This can be verified by unstacking the matrices in (2.7.4). See page 170 in chapter 7.

<sup>27</sup> 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.

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 savings program is concerned, activating a concern about robustness is equivalent with making the consumer more patient. However, that is not the end of the story. Chapter 13 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

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

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

$$k_t + c_t = Rk_{t-1} + d_t \quad (2.8.2a)$$

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

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

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

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

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

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

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

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

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

We provide a formula for the scalar  $\nu$  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  $\mu_{ct}$  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  $y'_t = [1 \quad k_{t-1} \quad d_t]'$  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

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



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  $R = 0_{3 \times 3}$  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  $\mu_{ct}$

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

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

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

Equations (2.8.8) and (2.8.9) imply that

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

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  $\nu$ .

#### 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  $\mu_{ct}$  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  $\epsilon_{t+1}$  is an i.i.d. scalar process with mean zero and unit variance.<sup>28</sup> 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.

---

<sup>28</sup> We computed  $\rho, c_d$  by calculating autocovariances implied by HST's specification, then used them to calculate the implied population first-order autoregressive representation.

- 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

$$\mu_{c,t+1} = \mu_{c,t} + 4.3825\epsilon_{t+1} \quad (2.8.12)$$

where we compute the coefficient 4.3825 on  $\epsilon_{t+1}$  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  $\sigma = \hat{\sigma} = -2E^{-7}$ .<sup>29</sup> We now compute that<sup>30</sup>

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

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  $\beta$  to  $\tilde{\beta} = .9995$ . Remarkably, with  $(\sigma, \beta) = (0, \tilde{\beta})$ , we compute that  $\mu_{c,t+1}$  obeys exactly (2.8.13).<sup>31</sup> Thus, starting from  $(\sigma, \beta) = (0, \hat{\beta})$ , insofar as the effects on consumption and saving are concerned, activating a concern about robustness by lowering  $\sigma$  while keeping  $\beta$  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, namely, 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  $\beta$ . In chapter 10, we extend these numerical examples analytically

<sup>29</sup> We discuss how to calibrate  $\sigma$  in chapters 9, 10, 13, and 14.

<sup>30</sup> 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.

<sup>31</sup> We discover this computationally using the method of the previous footnote.

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

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

Then the laws of motion for  $\mu_{c,t}$ , and therefore the decision rules for  $c_t$ , are identical across these two specifications of concerns about robustness. We establish formula (2.8.14) in appendix B of chapter 10.

### 2.8.5. Observational equivalence of quantities but not continuation values

We have seen that, holding other parameters constant, there exists a locus of  $(\sigma, \beta)$  pairs that imply the same consumption-savings 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  $(\sigma, \beta)$  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 13, 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 pairs imply identical savings and consumption plans, they imply different valuations of risky future consumption payoffs. In chapter 13, 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 36, 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 law<sup>32</sup>

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

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 motion of the state affects only the component  $d_t$  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 10, 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  $\mu_{ct}$  processes:

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

and

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

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

### 2.8.7. A Stackelberg formulation for representing misspecification

In chapters 7 and 8, we show the equivalence of outcomes under different timing protocols for the two-player zero-sum games. In appendix B of chapter 10, 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

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

<sup>32</sup> This is not a minimal state representation because we have not eliminated the constant from the  $Y$  component of the state.

subject to

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

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. 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. The optimal linear regulator has other uses too. In chapters 5, 17, and 18 we describe filtering problems. Via the concept of duality explained there, the linear regulator can also be used to solve such filtering problems, including those where the decision maker wants estimates that are robust to model misspecification.

Chapter 3 introduces a stochastic version of robust control problems and describes how they link to the non-stochastic problems of the present chapter. Chapters 4 and 5 then prepare the way for deeper studies of robust control and filtering problems by reviewing the foundations of ordinary (i.e., non-robust) control and filtering theory. In these two chapters, we shall encounter tools that will serve us well when we move on to construct robust decision rules and filters.

## A. Matlab programs

A robust optimal linear regulator is defined by the system matrices  $Q, R, A, B, C$ , the discount factor  $\beta$ , 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 more or less positive definite, where we say more or less because of the some detectability qualifications explained in chapter 4. 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 4 and by Hansen and Sargent (2008, 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.

## Chapter 3

### A stochastic formulation

*When Shannon had invented his quantity and consulted von Neumann on what to call it, von Neumann replied: ‘Call it entropy. It is already in use under that name and besides, it will give you a great edge in debates because nobody knows what entropy is anyway.’*

— Quoted by Georgii, “*Probabilistic Aspects of Entropy*,” 2003

#### 3.1. Introduction

This book makes ample use of the finding that the stochastic structure of a linear-quadratic-Gaussian robust control problem is a secondary concern because we can deduce robust decision rules by studying a related deterministic problem.<sup>1</sup> This chapter describes this finding in some detail. We start with a more general setting, then focus on the linear quadratic Gaussian case. We begin with a stochastic specification of shocks in an approximating model and describe misspecifications to that model in terms of perturbations to the distribution of the shocks. In the special linear-quadratic-Gaussian setting, formulas that solve the nonstochastic problem contain all of the information needed to solve a corresponding stochastic problem.<sup>2</sup>

#### 3.2. Shock distributions

Consider a sequence of i.i.d. Gaussian shocks  $\{\epsilon_t\}$  that enter the transition equation for an approximating model. The perturbed model alters the distribution of these shocks and, in particular, allows them to be temporally dependent. Let  $\epsilon^t = [\epsilon_t', \epsilon_{t-1}', \dots, \epsilon_1']'$ . Throughout, we will condition on the initial state  $y_0$ .<sup>3</sup> The date  $t$  information available to the decision maker is  $y_0$  and  $\epsilon^t$ .

---

<sup>1</sup> Some control theorists extend this insight beyond linear quadratic models and argue that stochastic structures are artificial and that all shocks should be regarded as deterministic processes that represent misspecifications. Although this interesting point of view has brought important insights, we don't embrace it. Instead, we strongly prefer to regard a model as a stochastic process and misspecifications as perturbations to a salient stochastic process that the decision maker takes as an approximating model.

<sup>2</sup> See Jacobson (1973).

<sup>3</sup> When some of the states are hidden from the decision maker, we would have to say more, as we do in Hansen and Sargent (2005b, 2007a) and in chapters 17 and 18. In this chapter, we will suppose that all of the state variables are observed.

### 3.3. Martingale representations of distortions

Following Hansen and Sargent (2005b, 2007a), we use martingales to represent distortions in the probabilities. This allows us to represent perturbed models by introducing some appropriately restricted multiplicative preference shocks into the original approximating model.

Let  $\pi(\varepsilon)$  denote the multivariate standardized normal distribution, where  $\varepsilon$  is a dummy variable with the same dimension as the number of entries of the random vector  $\varepsilon_t$ . Let  $\hat{\pi}(\varepsilon|\varepsilon^t, y_0)$  denote an alternative density for  $\varepsilon_{t+1}$  conditioned on date  $t$  information. Form the likelihood ratio

$$m_{t+1} = \frac{\hat{\pi}(\varepsilon_{t+1}|\varepsilon^t, y_0)}{\pi(\varepsilon_{t+1})}.$$

Notice that

$$E(m_{t+1}|\varepsilon^t, y_0) = \int \frac{\hat{\pi}(\varepsilon|\varepsilon^t, y_0)}{\pi(\varepsilon)} \pi(\varepsilon) d\varepsilon = 1,$$

where integration is with respect to the Lebesgue measure over the Euclidean space with the same dimension as the number of entries of  $\varepsilon_t$ . Now set  $M_0 = 1$  and recursively construct  $\{M_t\}$

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

Solving this recursion gives

$$M_t = \prod_{j=1}^t m_j.$$

The random variable  $M_t$  is a function of  $\varepsilon^t$  and  $y_0$  and evidently satisfies

$$E(M_{t+1}|\varepsilon^t, y_0) = M_t.$$

Hence,  $M_t$  is a martingale relative to the sequence of information sets (sigma algebras) generated by the shocks. The random variable  $M_t$  is a ratio of joint densities of  $\varepsilon^t$  conditioned on  $y_0$  and evaluated at the random vector  $\varepsilon^t$ . The numerator density  $\hat{\Pi}_t$  is the alternative one that we shall use to compute expectations.

Now let  $\phi(\varepsilon^t, y_0)$  be a random variable that is a Borel measurable function of  $\varepsilon^t$  and  $y_0$ , where  $\varepsilon^t$  is a dummy variable with the same dimension as the random vector  $\varepsilon^t$ . The expectation of  $\phi(\varepsilon^t, y_0)$  under the  $\hat{\pi}_t$  density can be computed as

$$\int \phi(\varepsilon^t, y_0) \hat{\Pi}_t(\varepsilon^t) d\varepsilon^t = E[M_t \phi(\varepsilon^t, y_0) | y_0]$$

where integration is with respect to the Lebesgue measure over a Euclidean space with the same dimension as the random vector  $\varepsilon^t$ .



### 3.4. A digression on entropy

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  $E(M_t \log M_t | y_0)$ . The function  $M_t \log M_t$  is convex in  $M_t$  and so lies above its linear approximation at the point  $M_t = 1$ . Thus,

$$M_t \log M_t \geq M_t - 1$$

because the derivative of  $M_t \log M_t$  is  $1 + \log M_t$  and equal to one for  $M_t = 1$ . Since  $E(M_t | y_0) = 1$ , it follows that

$$E(M_t \log M_t | y_0) \geq 0$$

and that  $E(M_t \log M_t | y_0) = 0$  only when  $M_t = 1$ , in which case there is no probability distortion associated with  $M_t$ .

The factorization  $M_t = \prod_{j=1}^t m_j$  implies the following decomposition of entropy:

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

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 (2007a) advocate using a discounted version of the object on the right side to penalize a malevolent player's choice of a sequence of increments  $\{m_{t+1}\}$ . Discounted entropy over an infinite horizon can be expressed as

$$(1 - \beta) \sum_{j=0}^{\infty} \beta^j E(M_j \log M_j | y_0) = \sum_{j=0}^{\infty} \beta^j E [M_j E(m_{j+1} \log m_{j+1} | \epsilon^j, y_0) | y_0],$$

where we have used a summation-by-parts formula. The right-hand side formula is particularly useful to us in recursive formulations of robust control problems in which we allow  $m_{t+1}$  to be chosen by a malevolent second agent at date  $t$ .

This formulation requires that the perturbed distributions  $\hat{\Pi}_t$  be *absolutely continuous* with respect to the baseline distribution  $\Pi_t$ . This means that the perturbed distribution cannot assign positive probability to events constructed in terms of  $\epsilon^t$  and  $y_0$  that have probability measure zero under the distribution implied by the approximating model.

### 3.5. A stochastic robust control problem

We want a robust decision rule for an action  $u_t$ . Suppose that the state evolves according to

$$y_{t+1} = \varpi(y_t, u_t, \epsilon_{t+1}),$$

the time period  $t$  return function is  $r(y_t, u_t)$ , and  $y_0$  is an initial condition. We require that a control process  $\{u_t\}$  have its time  $t$  component  $u_t$  be a function of  $\epsilon^t$  and  $y_0$ , so that  $y_t$  and  $r(y_t, u_t)$  also become functions of  $\epsilon^t$  and  $y_0$ .

To obtain a stochastic robust control problem that sustains an enduring concern about model misspecification, Hansen and Sargent (2007a) advocate using a two-player zero-sum game of the form

$$\max_{\{u_t\}} \min_{\{m_{t+1}\}} \sum_{t=0}^{\infty} E [\beta^t M_t \{r(y_t, u_t) + \alpha\beta E(m_{t+1} \log m_{t+1} | \epsilon^t, y_0)\} | y_0] \quad (3.5.1)$$

subject to

$$\begin{aligned} y_{t+1} &= \varpi(y_t, u_t, \epsilon_{t+1}), \\ M_{t+1} &= m_{t+1} M_t, \end{aligned} \quad (3.5.2)$$

where  $E m_{t+1} | \epsilon^t, y_0 = 1$ . Here  $\alpha \in [\underline{\alpha}, +\infty]$  is a penalty on the entropy associated with the  $m_{t+1}$  process. Soon we shall relate  $\alpha$  to  $\theta$  from chapter 2.

The two-person zero-sum game (3.5.1)-(3.5.2) has the player choosing processes for  $\{u_t\}_{t=0}^{\infty}$ ,  $\{m_{t+1}\}_{t=0}^{\infty}$  in a particular order. The dates on variables indicate informational constraints. We require that  $u_t$  be a function of  $\epsilon^t$  and  $y_0$  and  $m_{t+1}$  be a function of  $\epsilon^{t+1}$  and  $y_0$ .

### 3.6. A recursive formulation

Chapter 7 describes technical conditions that allow us to alter timing protocols without affecting outcomes and thereby to formulate an equivalent game that is recursive. The recursive game makes  $u_t$  a function of the Markov state  $y_t$  and  $m_{t+1}$  a function of  $\epsilon_{t+1}$  and the Markov state  $y_t$ , where  $m_{t+1}$  must have unit expectation.

To pose a recursive form of problem (3.5.1)-(3.5.2), we let the Markov state be the composite of  $M_t$  and  $y_t$ . We guess that the value function has the multiplicative form  $W(M, y) = MV(y)$  and consider the Bellman equation

$$\begin{aligned} MV(y) = \max_u \min_{m(\epsilon)} M \left\{ r(y, u) + \beta \int \left( m(\epsilon) V[\varpi(y, u, \epsilon)] \right. \right. \\ \left. \left. + \alpha m(\epsilon) \log m(\epsilon) \right) \pi(\epsilon) d\epsilon \right\} \end{aligned}$$

subject to the restriction that  $\int m(\varepsilon)\pi(\varepsilon)d\varepsilon = 1$ . The minimizing player chooses  $m$  as a function of  $\varepsilon$  so that  $m$  in effect is an infinite dimensional control vector.

The linear scaling of the value function by  $M$  allows us to consider the following problem that omits the state variable  $M$ :

$$V(y) = \max_u \min_{m(\varepsilon)} \left\{ r(y, u) + \beta \int \left( m(\varepsilon) V[\varpi(y, u, \varepsilon)] + \alpha m(\varepsilon) \log m(\varepsilon) \right) \pi(\varepsilon) d\varepsilon \right\} \quad (3.6.1)$$

subject to  $\int m(\varepsilon)\pi(\varepsilon)d\varepsilon = 1$ . A consequence of our being able to omit  $M_t$  as a state variable is that the control laws for  $u$  and  $m(\cdot)$  will depend on  $y$ , but not  $M$ .<sup>4</sup>

Consider the inner minimization problem

**Problem A:**

$$\mathcal{R}(V)(y, u) \equiv \min_{m(\varepsilon)} \int \left( m(\varepsilon) V[\varpi(y, u, \varepsilon)] + \alpha m(\varepsilon) \log m(\varepsilon) \right) \pi(\varepsilon) d\varepsilon$$

subject to  $\int m(\varepsilon)\pi(\varepsilon) = 1$ .

The objective is convex in  $m$  and the constraint is linear. The constraint  $\int m(\varepsilon)\pi(\varepsilon)d\varepsilon = 1$  restricts the average  $m(\cdot)$  but leaves open how to allocate  $m$  over alternative values of  $\varepsilon$ . Although  $m(\cdot)$  is infinite dimensional, it is easy to solve Problem A. Its solution is well known from the literature on relative entropy and large deviation theory.

The first-order conditions for minimization imply that

$$\log m(\varepsilon) = \frac{-V[\varpi(y, u, \varepsilon)]}{\alpha} + \lambda$$

where  $\lambda$  is Lagrange multiplier chosen so that  $\int m(\varepsilon)\pi(\varepsilon)d\varepsilon = 1$ . Therefore,

$$m^*(\varepsilon) = \frac{\exp\left(\frac{-V[\varpi(y, u, \varepsilon)]}{\alpha}\right)}{\int \exp\left(\frac{-V[\varpi(y, u, \varepsilon)]}{\alpha}\right) \pi(\varepsilon) d\varepsilon}. \quad (3.6.2)$$

Furthermore, under the minimizing  $m^*$ ,

$$\begin{aligned} & \int \left( m^*(\varepsilon) V[\varpi(y, u, \varepsilon)] + \alpha m^*(\varepsilon) \log m^*(\varepsilon) \right) \pi(\varepsilon) d\varepsilon \\ &= -\alpha \log \left[ \int \exp\left(\frac{-V[\varpi(y, u, \varepsilon)]}{\alpha}\right) \pi(\varepsilon) d\varepsilon \right] \\ &= \mathcal{R}(V)(y, u). \end{aligned}$$

This is the risk-sensitive recursion of Hansen and Sargent (1995).

---

<sup>4</sup> Our decision to use entropy to measure model discrepancies facilitates this outcome.

### 3.6.1. Verifying the solution

As a check on this calculation, write

$$\begin{aligned} \int m \log m \pi d\varepsilon &= \int \frac{m}{m^*} (\log m - \log m^*) m^* \pi d\varepsilon + \int m \log m^* \pi d\varepsilon \\ &\geq \int m \log m^* \pi d\varepsilon. \end{aligned}$$

This inequality follows because the quantity

$$\int \frac{m}{m^*} (\log m - \log m^*) m^* \pi d\varepsilon$$

is a measure of the entropy of  $m$  relative to  $m^*$  and hence is nonnegative. Thus,

$$\begin{aligned} \int mV[\varpi(y, u, \cdot)] \pi d\varepsilon + \alpha \int m \log m \pi d\varepsilon &\geq \\ \int mV[\varpi(y, u, \cdot)] \pi d\varepsilon - \int mV[\varpi(y, u, \cdot)] \pi d\varepsilon + \mathcal{R}(V)(y, u) & \\ = \mathcal{R}(V)(y, u), & \end{aligned}$$

where we have substituted using formula (3.6.2) for  $m^*$ . This verifies that  $m^*$  is the minimizer in Problem A.

## 3.7. A value function bound

The random function  $m^*$  of  $\varepsilon$  tilts the density of the shock  $\varepsilon$  exponentially using the value function to determine the directions where the decision maker is most vulnerable. Since  $m^*$  depends on the state  $y_t$ , the resulting distorted density for the shocks can make the shocks temporally dependent and thereby represent misspecified dynamics. The form of the worst-case density depends on both the original density  $\pi$  and the shape of the value function  $V$ . When  $\pi$  is normal and  $V$  is quadratic, the distorted density is normal.


As a direct implication of Problem A, we obtain a bound on the distorted expectation of the value function as a function of relative entropy:

$$\int mV[\varpi(y, u, \cdot)] \pi(\varepsilon) d\varepsilon \geq \mathcal{R}(V)(y, u) - \alpha \int m \log m \pi(\varepsilon) d\varepsilon. \quad (3.7.1)$$

The first term on the right depends on  $\alpha$  but not on the alternative model as characterized by  $m$ . The second term is  $-\alpha$  times entropy. Thus, inequality (3.7.1) justifies interpreting  $\alpha$  as a *utility price of robustness*. The larger is the relative entropy, the larger is the downward adjustment in the relative entropy bound.

### 3.8. Large deviation interpretation of $\mathcal{R}$

We have interpreted Problem A in terms of a concern about robustness that is achieved by substituting the operator  $\mathcal{R}$  for the conditional expectations operator in a corresponding Bellman equation without a concern for robustness. Let  $y^+$  denote the state next period. In this section, we use ideas from the theory of large deviations to indicate how the operator  $\mathcal{R}(V)(y, u)$  contains information about the left tail of the distribution of the continuation value  $V(y^+)$  where the distribution of  $y^+ = \varpi(y, u, \varepsilon)$  is induced by the density  $\pi(\varepsilon)$  associated with the approximating model. Recall from Problem A that  $\mathcal{R}$  depends on  $\alpha$  and collapses to the conditional expectation operator as  $\alpha \nearrow +\infty$ . We shall show that  $\mathcal{R}$  contains more information about the left tail of  $V$  as  $\alpha$  is decreased. We gather this interpretation from an exponential inequality that bounds the (conditional) tail probabilities of the continuation value. This tail probability bound shows 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.



**Figure 3.8.1:** 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  $\alpha$ : 1 and 2.

The tail probability bound is widely used in the theory of large deviation approximations.<sup>5</sup> It uses the inequality

$$\mathbf{1}_{\{V:V \leq -r\}} \leq \exp\left[\frac{-(V+r)}{\alpha}\right]$$

<sup>5</sup> For an informative survey, see Bucklew (1990).

depicted in figure 3.8.1, where  $\mathbf{1}$  is the indicator function. This inequality holds for any real number  $r$  and any  $\alpha > 0$ . Then computing expectations conditioned on the current state vector  $y$  and control  $u$  yields

$$\text{Prob}\{V(y^+) \leq -r|y, u\} \leq E \left( \exp \left[ -\frac{V(y^+)}{\alpha} \right] \middle| y, u \right) \exp \left( -\frac{r}{\alpha} \right)$$

or

$$\text{Prob}\{V(y^+) \leq -r|y, u\} \leq \exp \left[ -\frac{\mathcal{R}(V|y, u)}{\alpha} \right] \exp \left( -\frac{r}{\alpha} \right). \quad (3.8.1)$$

Inequality (3.8.1) bounds the tail probability on the left by an exponential in  $r$ . Thus,  $\alpha$  determines a decay rate in the tail probabilities of the continuation value. Decreasing  $\alpha$  increases the exponential rate at which the bound sends the tail probabilities to zero, thereby expressing how a lower  $\alpha$  heightens concern about tail events. Associated with this rate is a scale factor

$$\int \exp \left( -\frac{V[\varpi(y, u, \varepsilon)]}{\alpha} \right) \pi(y, u, \varepsilon) d\varepsilon = \exp \left[ -\frac{\mathcal{R}(V|y, u)}{\alpha} \right].$$

The adjustment to the value function determines the constant associated with the prespecified decay rate. For a fixed  $\alpha$ , a larger value of  $\mathcal{R}(V)(y, u)$  gives a smaller scale factor in the probability bound.

### 3.9. Choosing the control law

To construct a robust control law, solve the outer maximization problem of (3.6.1)

$$\max_u r(y, u) + \beta \mathcal{R}(V)(y, u).$$

Notice that we computed  $m$  as a function  $(y, u)$  before solving for  $u$ . It is often the case that we could compute  $m$  and  $u$  simultaneously as functions of  $y$  by in effect stacking first-order conditions instead of proceeding in sequence. This justifies an algorithm for the linear quadratic case that we describe in section 2.7 of chapter 2.

### 3.10. Linear-quadratic model

To connect the approach of this chapter to the nonstochastic formulations summarized in chapter 2, we turn to a linear quadratic setting with Gaussian disturbances. Consider the following evolution equation:

$$y^+ = Ay + Bu + C\varepsilon$$

where  $y^+$  is the next period value of state vector. Consider a value function

$$V(y) = -\frac{1}{2}y'Py - \rho.$$

From our previous calculations, we know that

$$m^*(\varepsilon) \propto \exp \left[ \frac{1}{2\alpha} \varepsilon' C' PC \varepsilon + \frac{1}{\alpha} \varepsilon' C' P (Ay + Bu) \right].$$

When  $\pi$  is a standard normal density, it follows that

$$\begin{aligned} \pi(\varepsilon) m^*(\varepsilon) \propto \exp \left[ -\frac{1}{2} \varepsilon' \left( I - \frac{1}{\alpha} C' PC \right) \varepsilon \right. \\ \left. + \varepsilon' \left( I - \frac{1}{\alpha} C' PC \right) (\alpha I - C' PC)^{-1} C' P (Ay + Bu) \right], \end{aligned}$$

where the proportionality coefficient is chosen so that the function of  $\varepsilon$  on the right-hand side integrates to unity. The right-hand side function can be recognized as being proportional to a normal density with covariance matrix  $(I - \frac{1}{\alpha} C' PC)^{-1}$  and mean  $(\alpha I - C' PC)^{-1} C' P (Ay + Bu)$ . Evidently, the covariance matrix of the shock is enlarged. The altered mean for the shock implies that the distorted conditional mean for  $y^+$  is

$$\left[ I + C (\alpha I - C' PC)^{-1} C' P \right] (Ay + Bu).$$

These formulas for the distorted means of  $\varepsilon$  and  $y^+$  agree with formulas that we derived from a deterministic problem in chapter 2.

### 3.11. Relative entropy and normal distributions

As we have just seen, the worst-case distribution will also be normal. As a consequence, we consider the corresponding measure of relative entropy for a normal distribution. This renders the following calculation interesting. Suppose that  $\pi$  is a multivariate standard normal distribution and that  $\hat{\pi}$  is normal with mean  $w$  and nonsingular covariance  $\Sigma$ . We seek a formula for  $\int (\log \hat{\pi}(\varepsilon) - \log \pi(\varepsilon)) \hat{\pi}(\varepsilon) d\varepsilon$ . First, note that the likelihood ratio is

$$\log \hat{\pi}(\varepsilon) - \log \pi(\varepsilon) = \frac{1}{2} \left[ -(\varepsilon - w)' \Sigma^{-1} (\varepsilon - w) + \varepsilon' \varepsilon - \log \det \Sigma \right] \quad (3.11.1)$$

To compute relative entropy, we must evaluate expectations using a normal distribution with mean  $w$  and covariance  $\Sigma$ . Observe that

$$- \int \frac{1}{2} (\varepsilon - w)' \Sigma^{-1} (\varepsilon - w) \hat{\pi}(\varepsilon) d\varepsilon = \frac{1}{2} \text{trace}(I).$$

Applying the identity  $\varepsilon = w + (\varepsilon - w)$  gives

$$\frac{1}{2}\varepsilon'\varepsilon = \frac{1}{2}w'w + \frac{1}{2}(\varepsilon - w)'(\varepsilon - w) + w'(\varepsilon - w).$$

Taking expectations,

$$\frac{1}{2}\int \varepsilon'\varepsilon\hat{\pi}(\varepsilon) d\varepsilon = \frac{1}{2}w'w + \frac{1}{2}\text{trace}(\Sigma).$$

Combining terms gives

$$\int (\log \hat{\pi} - \log \pi) \hat{\pi} d\varepsilon = -\frac{1}{2}\log \det \Sigma + \frac{1}{2}w'w + \frac{1}{2}\text{trace}(\Sigma - I). \quad (3.11.2)$$

### 3.12. Value function adjustment for the LQ model

Our adjustment to the value function is

$$\mathcal{R}(V)(y, u) = -\alpha \log \left[ \int \exp \left( \frac{-V[\varpi(y, u, \varepsilon)]}{\alpha} \right) \pi(\varepsilon) d\varepsilon \right].$$

For linear quadratic problems, we have at our disposal a more explicit depiction of this adjustment. Recall that this adjustment is given by  $\int V\hat{\pi}d\varepsilon + \alpha \int (\log \hat{\pi} - \log \pi)\hat{\pi}d\varepsilon$  for the  $\hat{\pi}$  obtained as the solution to the minimization problem. As we have already shown,  $\hat{\pi}$  is a normal density with mean  $(\alpha I - C'PC)^{-1}C'P(Ay + Bu)$  and covariance matrix  $(I - \frac{1}{\alpha}C'PC)^{-1}$ . Using our earlier calculations of relative entropy (3.11.2), the adjustment to the linear-quadratic objective function,  $-\frac{1}{2}y'Py - \rho$  is<sup>6</sup>

$$\begin{aligned} \mathcal{R}(V)(y, u) = & -\frac{1}{2}(Ay + Bu)' \left[ P + PC(\alpha I - C'PC)^{-1}C'P \right] (Ay + Bu) - \rho \\ & + \frac{\alpha}{2}\text{trace} \left[ \left( I - \frac{1}{\alpha}C'PC \right)^{-1} - I \right] \\ & - \frac{\alpha}{2}\log \det \left( I - \frac{1}{\alpha}C'PC \right)^{-1}. \end{aligned}$$

It is enough to work with a deterministic counterpart to this adjustment in the linear quadratic case. For the purposes of computation, consider the following deterministic evolution for the state vector:

$$y^+ = Ay + Bu + Cw$$

---

<sup>6</sup> This expression motivates setting  $\theta$  in chapter 2 equal to  $\alpha/2$  in order to match up with the formulation in this chapter.



where we have replaced the stochastic shock by a distorted mean  $w$ . Since this is a deterministic evolution, covariance matrices do not come in play now. Solve the problem

$$\min_w -\frac{1}{2} (Ay + Bu + Cw)' P (Ay + Bu + Cw) + \frac{\alpha}{2} w' w.$$

In this problem, relative entropy is no longer well defined. Instead, we penalize the choice of the distortion  $w$  using only the contribution to relative entropy (3.11.2) coming from the mean distortion. The solution for  $w$  is

$$w^* = (\alpha I - C' P C)^{-1} C' P (Ay + Bu).$$

This coincides with the mean distortion of the worst-case normal distribution described earlier. The minimized objective function is

$$-\frac{1}{2} (Ay + Bu)' \left[ P + P C (\alpha I - C' P C)^{-1} C' P \right] (Ay + Bu),$$

which agrees with the contribution to the stochastic robust adjustment to the value function coming from the quadratic form in  $(Ay + Bu)$ . What is missing relative to the stochastic problem is the distorted covariance matrix for the worst-case normal distribution and the constant term in the adjusted value function. However, neither of these objects alters the computation of the robust decision rule for  $u$  as a function of the state vector  $y$ .

This trick underlies much of the analysis in the book. For the purposes of computing and characterizing the decision rules in the linear-quadratic model, we can focus exclusively on mean distortions and can abstract from covariance distortions. In the linear-quadratic case, the covariance distortion alters the value function only through the additive constant term. Using and refining the formulas in this chapter, we can deduce both the covariance matrix distortion and the constant adjustment. As we shall see, these ideas also apply when we turn to issues involving decentralization and welfare analysis.