

# **Robustness, Detection and the Price of Risk**

**Evan W. Anderson**

University of North Carolina

**Lars Peter Hansen**

University of Chicago

**Thomas J. Sargent**

Stanford University and Hoover Institution

*March 27, 2000*

Note: We thank Jose Scheinkman, Eric Renault, Grace Tsiang, and Neng Eric Wang for comments on an earlier draft. This is a revision of an previous paper entitled *Risk and Robustness in Equilibrium*.

## 1. Introduction

A model is a probability distribution over a sequence. Rational expectations models assume that agents know the model and are not concerned about specification error. Agents' confidence in the model specification contrasts sharply with the attitudes of both econometricians and calibrators.

Econometricians routinely use likelihood-based specification tests (information criteria or IC) to select from a set of models. Those tests are used to organize comparisons of empirical distributions with plausible models. Less formally, calibrators sometimes justify their estimation procedures by saying that they regard their models as incorrect. Nevertheless, the agents inside a calibrator's model do not share the model-builders' doubts about specification. Thus, the rational expectations econometrician or calibrator typically attributes no concern about specification error to agents even as he/she shuttles among alternative specifications.

This paper is about models with agents whose doubts about model specification cause them to value decision rules that perform well across a set of models. Agents fear difficult-to-detect misspecifications of the state transition law, difficult to detect because they are partly masked by the random shocks that impinge on the dynamical system. In response, agents adjust decision rules to guard against the modeling errors. This makes more cautious decisions and puts model uncertainty premia into equilibrium security market prices.

We extend earlier work of Hansen and Sargent (1995) and Hansen, Sargent, and Tallarini (HST, 1999) in several ways. First, we go beyond the linear-quadratic framework of Hansen and Sargent (1995) and HST, by showing how to compute robust decision rules and prices in economies with general return and transition functions and a representative agent who prefers a robust rule. In both discrete and continuous time, we produce economical and manageable generalizations of the Bellman and asset pricing equations associated with a rational expectations model. The generalizations modify the expectation operators in those equations, twisting the probability distribution associated with a rational

---

<sup>1</sup> This assumption is so widely used that it rarely excites comment within macroeconomics. Kurz is an exception. The rational expectations critique of earlier dynamic models with adaptive expectations was that they implicitly contained two models: one for the econometrician; a worse one for the agents inside the model doing forecasting. See Lucas (1976) and Jorgenson (1967) for statements of this critique in alternative contexts. Rational expectations modeling responded to this critique by attributing a common model to the econometrician and the agents within his model. The econometricians and the agents can have different information sets, but they agree about the model (stochastic process).

<sup>2</sup> Which after all are likelihood functions.

<sup>3</sup> For example, see the two papers about specification error in rational expectations models by Sims (1993) and Hansen and Sargent (1993).

expectations model. The twisting is governed by a single parameter measuring the set of alternative specifications that concern the decision maker.

Second, we explore the effect of this parameter on asset prices through the market price of risk. We show how concern for model misspecification adds components to what are usually referred to as the factor risk prices. As a consequence, the implied risk-return tradeoff observed from security market data has a model uncertainty component that we call the market price of uncertainty. We link the market price of uncertainty to a model-discrimination measure, Chernoff entropy, that governs the probabilities of distinguishing alternative models from time-series data. For a given data set, the dynamic measure of Chernoff entropy locates specifications that statistical tests cannot decide among at reasonable confidence levels. We use the theory of statistical detection to ascertain plausible preferences for robustness and magnitudes for the market price of uncertainty.

### 1.1. *Related literature*

In the context of a discrete-time, linear-quadratic permanent income model, Hansen, Sargent, and Tallarini (HST) consider a particular class of model misspecifications that are restrained in terms of a single robustness parameter. HST showed how robust decision-making induces behavior similar to that produced by risk aversion. They interpreted a preference for robustness as aversion to Knightian uncertainty. They calculated how much preference for robustness is required to put market prices of risk into empirically plausible regions.

Hansen and Sargent (1999) and HST allowed a limited array of specification errors that take the form of shifts in the conditional mean of shocks that would be i.i.d. and normally distributed under the approximating model. In this paper, we consider more general environments and motivate the form of specification errors more directly in terms of specification test statistics. We can show that the particular perturbations assumed by HST then emerge in the special linear-quadratic, Gaussian control problems and also in a more general class of control problems in which the stochastic evolution of the state is a Markov diffusion process. Specifications different from HST's are required when Markov jump processes are entertained. As in HST, our formulation of robustness allows us to

---

<sup>4</sup> Econometricians often use specification test statistics to classify alternative possible specifications into two sets, those that the data don't distinguish themselves from, and those rejected specifications that seem remote from the data. Depending on the confidence level set for the test and the data set, the former set is a collection that the IC leave as candidate specifications. The econometrician typically emerges from an empirical specification analysis with a preferred specification and at least implicitly a penumbra of alternative specifications that are statistically close to the preferred one. We propose to infuse such specification doubts into the agents.

reinterpret one of Epstein-Zin's (1989) recursions as reflecting a preference for robustness rather than aversion to risk.

### *1.2. Robustness versus learning*

Our agents make decisions using an approximating model. To quantify the approximation, we measure discrepancy between the approximating model and other models in terms of relative entropy, an expected log likelihood ratio, where the expectation is taken with respect to the distribution from the alternative model. Relative entropy is used in the theory of large deviations, a powerful mathematical theory about the rate at which uncertainty about unknown distributions is resolved as the number of observations grows. Since we use an entropy concept to restrain model perturbations, we appeal to the theory of statistical detection to provide information about a preference for robustness that is quantitatively reasonable. We make our decision-makers prefer robustness against alternatives that are difficult to detect statistically.

The perspective of a robust decision maker differs substantially from that of one who learns. In our dynamic settings, the robust decision maker accepts model misspecification as a permanent state of affairs, and devotes his thoughts to designing robust controls, rather than to using data to improve his model specification over time. The robust decision-maker turns his back on learning. By way of contrast, many formulations of learning have decision-makers fully embrace an approximate model when making their choices. Despite their different orientations, robust decision makers and ones who learn share a need for a convenient way of measuring and interpreting the proximity of two probability distributions. This need places technical bridges between robust decision theory and learning theory. In learning theory, expressions from large deviation theory provide bounds on rates of learning. Those same expressions provide bounds on value functions across alternative possible models in robust decision theory.

---

<sup>5</sup> See Bray (1982) and Kreps (1998).

### 1.3. *Organization*

The remainder of this paper assembles and applies tools for analyzing and quantifying the existence of situations in which a decision maker would be wise to value robustness because he suspects that the data are governed by one of a set of models that a ‘learner’ could not confidently choose among after many observations. Section 2 gives an overview of the paper by altering a planning problem that underlies the standard asset pricing model used in finance and macroeconomics. Section 3 describes an inequality from the theory of large deviations and uses it to interpret a twisted expectation operator as representing a preference for robustness. Section 4 introduces relative entropy as a measure of discrepancy between models and uses it to represent a value function under a concern for robustness in discrete time. Section 5 formulates things in continuous time, and section 6 specializes to the case of diffusions. Section 7 studies equilibrium pricing under robustness and derives an additive decomposition of the prices of Brownian motion increments in terms of a risk component and a model ambiguity component. Section 8 uses a stylized model selection problem to deduce an instantaneous statistical distance between two models. This distance results in a Chernoff measure of entropy that, in the case of diffusions, is essentially equivalent to the relative entropy measure we use to restrain robustness. This link allows us to use the theory of statistical discrimination to help calibrate the robustness parameter. Section 9 gives a formula for the implied risk-return tradeoff and relates the slope of the tradeoff to the entropy measure of statistical discrepancy between two models. It then considers two examples linking the market price of risk to detection probabilities. Section 10 states our conclusions and plans.

## 2. **Overview**

To illustrate a preference for robustness within an equilibrium model, we will consider a model familiar to both macroeconomists and financial economists. Following Lucas and Prescott (1971) and Brock and Mirman (1972), we begin with a planning problem. The solution provides the intertemporal allocation of consumption, investment, and capital. The shadow prices of the planning problem can be used to construct competitive equilibrium state-date prices.

### 2.1. Robust Resource Allocation

A planner has the approximating model

$$dx_t = \mu(x_t, i_t) dt + \Lambda(x_t, i_t) dB_t \quad (2.1)$$

where  $x_t$  is a vector of state variables including endogenous states such as capital stocks and exogenous states used to model the persistence in the underlying risk,  $i_t$  is a vector of control variables, including say investment, and  $\{B_t\}$  is a vector standard Brownian motion. The time  $t$  drift (local mean) vector is  $\mu(x_t, i_t)$ . Both the drift and the diffusion (local covariance) matrix  $\Sigma(x_t, i_t) = \Lambda(x_t, i_t)' \Lambda(x_t, i_t)$  can depend only on the control and the state. The control is limited to depend only on past values of the Brownian motion vector. The objective of the resource allocation is to maximize

$$E \int_0^\infty \exp(-\delta t) U(x_t, i_t) dt, \quad (2.2)$$

subject to (2.1) where  $U(x_t, i_t)$  measures the instantaneous utility,  $\delta$  is a subjective rate of discount. We are interested in models for which this optimal resource allocation problem has a Markov solution of the form  $i_t = f(x_t)$ .

The planner regards his model as an approximation to an unknown true model. The planner thinks that the true model appends an unknown drift to the Brownian motion in his approximating model (2.1), so that in the true model the Brownian motion  $B_t$  in (2.1) is replaced by

$$B_t + \int_0^t g_s ds.$$

The process  $\{g_t\}$  must be adapted to the Brownian motion filtration but otherwise can depend arbitrarily on the history of the state vector up to  $t$ . Thus,  $g_t$  captures forms of model misspecification that are not Markovian. The true state evolution equation is:

$$dx_t = \mu(x_t, i_t) dt + \Lambda(x_t, i_t) (g_t + dB_t) \quad (2.3)$$

We shall call  $g_t$  the ‘distortion’ of the true model relative to the approximating model. Later we shall justify in detail why the distortion takes this form. For now, we note that the distortion  $g_t$  could represent misspecifications of the exogenous dynamics, or of the mechanism by which the control  $i_t$  alters the endogenous capital stocks, provided this mechanism is stochastic. By construction the drift is at least partially masked by the Brownian motion. We will be interested in restraining  $\{g_t\}$  in a way to represent forms of model misspecification that are difficult to detect statistically.

To make the decision problem well posed, we must attribute a view about the distortion  $g_t$  to the decision maker. One way to proceed would be to give the decision maker a prior

distribution over a class of  $g_t$  processes. However, we do not let the decision-maker have such a precise probabilistic description of the potential forms of model misspecification: in our view, that would give him the correct specification. It would eliminate the fear of model misspecification and any notion that the final model is an approximation.

We have in mind a situation where the decision maker can't form a prior, motivated in part by Knight (1921) and Ellsberg (1971). The potential misspecifications are on the one hand complicated, but on the other hand difficult to distinguish statistically from our original model.

EXAMPLE 2.1. *Knight's Urns.* Knight (1921) distinguished between risk and uncertainty, reserving the latter term for events for which objective probabilities could not plausibly be assigned. Savage (1972) axiomatized a personal or Bayesian theory of decision making that undermined Knight's distinction. For Savage, the source of the personal probabilities underlying agents' decisions was beyond question and irrelevant. But Ellsberg (1971) created an example that challenged the Bayesian-Savage model of decision-making. Ellsberg (1971) considered a choice between bets on two urns. In Urn A, it is known in advance that there are fifty red balls and fifty black balls. In Urn B, the fractions of red and black balls are not known in advance. A ball will be drawn randomly from each urn. At no cost a decision maker is permitted to guess the color of the ball drawn from one and only one of the urns. If the decision maker guesses the color drawn from the chosen urn, he receives a positive payoff. The Bayesian-Savage model predicts either indifference between urns or a preference for Urn B, depending on the prior probability assignment. But Ellsberg (1971) argued that a preference for Urn A is reasonable. To Ellsberg, there is an important distinction between urns for which you are informed about probabilities *vis-a-vis* ones for which you are not. Gilboa and Schmeidler (1989) use this urn example to motivate a relaxation of the Bayesian model in which multiple priors may justify decisions under a *max-min* criterion. The ambiguity about the assignment of priors is a way to make Knight's notion of uncertainty operational. An aversion to such uncertainty can rationalize a preference for Urn A. When the decision-maker assigns a nondegenerate range of beliefs about the

---

<sup>6</sup> For example, the model misspecification might take the form of a hidden state variable shifting the investment opportunities. By solving the combined filtering and control problem, we could reinterpret it as a resource allocation problem with additional state variables used to keep track of posterior probabilities of the hidden Markov states (*e.g.* see Elliot, Aggoun and Moore, 1995).

<sup>7</sup> Epstein and Wang (1994) also cite the Ellsberg paradox for motivating their dynamic formulation of a multiple priors model.

<sup>8</sup> As Ellsberg (1971) emphasized, Ellsberg and Knight drew some rather different conclusions using this hypothetical gamble.

<sup>9</sup> For instance, if the decision maker believes there are more red balls than black balls in Urn B (assigns a prior probability of drawing a red ball that is greater than one-half), then he strictly prefers betting on Urn B.

fraction of red balls in Urn B that includes .5, a preference for Urn A can be justified on the basis of an aversion to Knightian uncertainty.

The potential distortions in the state evolution indexed by  $\{g_t\}$  give rise to a family of models of the state evolution. This family plays the role of the multiple priors of Gilboa and Schmeidler (1989). Our next example illustrates some complicated distortions in a familiar model of macroeconomics and finance.

EXAMPLE 2.2. *A Stochastic Growth Model.* We consider a version of a model due to Brock (1979), Cox, Ingersoll and Ross (1985), and Brock and Magill (1979). There are multiple technologies for transferring goods from one instant to the next. Capital is freely transferable across the technologies. Newly produced output is split between consumption and new capital. Let  $k_{it}$  denote the capital stock allocated to technology  $i$ , and let  $k_t$  denote aggregate capital. Then aggregate capital evolves according to

$$dk_t = \sum_i [(k_{it})^{\alpha_i} dy_{it} - \delta_i k_{it} dt] - c_t dt$$

where

$$dy_{it} = \mu_i(z_t) dt + \sigma_i(z_t) \cdot (g_t + dB_t).$$

The  $\alpha_i$ 's, the capital productivity parameters, are in the interval  $(0, 1]$  and the  $\delta_i$ 's are the depreciation rates for capital. The date  $t$  exogenous state vector  $z_t$  can shift both the local mean and the local variance of the change in technology  $dy_{it}$ . The state vector  $z_t$  evolves according to

$$dz_t = \mu_z(z_t) dt + \Lambda_z(z_t) (g_t + dB_t).$$

Here  $z_t$  is a vector,  $\mu_z$  is a vector, and  $\Lambda$  is a matrix with the appropriate dimensions. Here  $g_t \neq 0$  represents the gap between the approximating model (for which  $g_t \equiv 0$ ) and the unknown true model.

The control vector is consumption and the fraction of aggregate capital allocated to each technology. Limiting the fractions to be nonnegative will enforce a nonnegativity constraint on the individual capital stocks, and since the fractions add to unity, the dimension of the control vector coincides with the number of technologies.

This model is closely related to many in the literature. When there is a single technology, it is a continuous-time version of a Brock-Mirman (1972) economy. As noted by Brock (1979), when the  $\alpha_i$ 's are all unity this resource allocation problem collapses to a version of

---

<sup>10</sup> While there is no counterpart to the  $\{z_t\}$  process in Brock's (1979) model, there is in the version due to Cox, Ingersoll and Ross (1985).

Merton's (1973) *portfolio* allocation problem posed instead as a resource allocation problem. It is also a version of the Cox, Ingersoll and Ross's (1985) general equilibrium model of financial markets.

The perturbations  $g_t$  in the Brownian motion can represent misspecifications of the exogenous process  $\{z_t\}$  shifting the technology opportunities or misspecifications in the productivity of capital, provided that this productivity is disguised by a Brownian motion. Since the marginal productivity is modeled as  $\alpha_i(k_{it})^{1-\alpha_i}[\mu_i(z_t)dt + \sigma_i(z_t) \cdot (g_t + dB_t)] - \delta_i$ , a nondegenerate  $\sigma_i$  is needed to conceal such a distortion in the Brownian motion drift.

The decision maker wants a decision rule that will work well across a set of  $g_t$ 's that are close to zero. To promote robust decision-making, we replace the single agent planner by two decision makers with conflicting aims. We alter the instantaneous objective in (2.2) by appending a penalty term  $\frac{\theta}{2}g_t'g_t$  so that the intertemporal criterion becomes

$$E \int_0^\infty \exp(-\delta t) \left[ U(x_t, i_t) + \frac{\theta}{2}g_t'g_t \right]. \quad (2.4)$$

The penalty term is introduced to restrain the choice of  $g_t$ . To deduce the equilibrium quantity allocation, we find the Nash Markov equilibrium of a two-player game where one player minimizes (2.4) by choice of control  $\{g_t\}$  and the other maximizes (2.4) by choice of the  $\{i_t\}$ .

The macroeconomic literature on representative agent, dynamic, stochastic equilibrium models under rational expectations typically considers models (specified in discrete time) in which  $g_t = 0$  (enforced by setting  $\theta = \infty$ ). Setting  $\theta < +\infty$  lets  $g_t \neq 0$  and fear of model misspecification into the analysis. For sufficiently small values of  $\theta$  there may fail to exist equilibria of the two-player game. It is convenient and restrictive that we penalize large deviations  $g_t$  with a single parameter  $\theta$ , say, as opposed to say constraining  $g_t$  instant by instant.

Choosing  $\{g_t\}$  in the malevolent way embodied in the two person game represents a preference for a decision rule for  $\{i_t\}$  that is robust to a wide variety of model misspecifications. The parameter  $\theta$  indexes the amount of robustness sought. A larger value of  $\theta$  more strictly restrains the malevolent player and weakens the incentives to be robust. In what follows, we let  $\{i_t^*\}$  and  $\{g_t^*\}$  denote the state-contingent decisions of the two players.

---

<sup>11</sup> When there are multiple capital stocks, the control vector influences the  $\Lambda$  that enters the composite state vector process. In James's (1992) continuous-time development of the link between robustness and risk sensitivity, this dependence is absent. This influence is also abstracted from in the linear-quadratic formulation of the robust control problem.

<sup>12</sup> The robustness penalty parameter  $\theta$  can typically be interpreted as a Lagrange multiplier on a discounted intertemporal specification-error constraint. See Hansen and Sargent (2000) for more discussion.

The worst-case process  $\{g_t^*\}$  should be compared to two other versions of  $g_t$ . We represent the decision maker's approximating model by setting  $g_t \equiv 0$ . The *true* model is represented in terms of some unknown  $\hat{g}_t \neq 0$ . In general,  $g_t^* \neq \hat{g}_t$ . The worst case process  $\{g_t^*\}$  is not an estimate of  $\hat{g}_t$ , but rather an instrument for designing a rule that is robust against a set of possible misspecifications.

We shall formulate the min-max game associated with (2.4) recursively. More precisely, we adopt timing and information protocols that lead us to focus on the Markov-perfect equilibrium of the game, in which the decision maker and the malevolent opponent choose state-feedback rules for  $i_t$  and  $g_t$ , respectively. This shall lead us to formulate the zero-sum game in terms of a single Bellman equation. With these timing protocols, at any date neither of the players commits to a future sequence of decisions. The standard argument for imposing subgame perfection has special appeal in the current setting, because the discrepancy  $\hat{g}_t - g_t^*$  means that the agent understands that the actual state  $x_t$  will certainly wander off the equilibrium path of the zero-sum game. Thus, the agent concedes that the actual law of motion is

$$dx_t = \mu[x_t, i^*(x_t)] dt + \Lambda[x_t, i^*(x_t)] [\hat{g}_t dt + dB_t]$$

for some unknown  $\{\hat{g}_t\}$  process.

## 2.2. Statistical Detection

In the Urn Example 2.1, the decision maker could benefit from repeated draws from Urn B. These draws would have him learn more about the number of red and black balls with repeated samples from the same urn. This information might eventually reveal that it is better to bet on Urn B. If, however, the repeated data are draws from the different urns and changes in the urn are not observed, then Dow and Werlang (1994) characterize cases in which the Knightian uncertainty will remain. If the resulting interval of probabilities contains one-half, then the strict preference for Urn A can still be defended.

In our analysis, the robust planner does not attempt to learn more about the model misspecification. As in any rational expectations model, some prior learning process is embodied in the selection of the benchmark model to be used in decision making. Our choice of  $\theta$  will restrain the amount of model misspecification in a particular statistical sense. The idea is that we do want the planner to guard against alternatives that are difficult to detect from historical data. We could follow Dow and Werlang (1994) and study model

---

<sup>13</sup> In the control theory literature, what we call the approximating model is often called the reference model.

perturbations that are *impossible* to detect. Instead, we appeal to a statistical discrimination theory originated by Chernoff to measure model misspecification. In disguising model misspecification by Brownian motion, we can make statistical detection *difficult*.

To formalize this difficulty, we use the continuous-time Markov process extension due to Newman and Stuck (1979) to justify the measure  $\frac{g_t' g_t}{8}$  as the detection error rate between competing models. By deducing how  $\theta$  restrains the detection error rate  $\frac{g_t^* g_t^*}{8}$ , we will set  $\theta$  to exclude model misspecifications that are easily detectable using time series data on the Markov state.

### 2.3. Factor pricing

Following Lucas (1978) and Breeden (1979), we extract equilibrium security market prices as shadow prices for the planner. In particular, we imitate Breeden (1979) by constructing factor risk prices of the Brownian motion increments that depict a risk-return tradeoff. We enrich the tradeoff by adding prices of model uncertainty that are induced by the planner's preference for robustness. Let  $\mu_r(x_t)$  denote the instantaneous rate of return on a security with a time  $t$  Brownian increment  $\beta(x_t) \cdot dB_t$ . Then the Euler equation pricing formula restricts  $(\mu_r, \beta)$  via

$$\mu_r - \rho = -\beta \cdot (\sigma_m + g^*)$$

where  $\rho$  is the (instantaneous) riskfree rate,  $\sigma_{m,t}$  is a vector of factor risk prices and  $g_t^*$  is a vector of model uncertainty prices. The factor risk prices have the same form as Breeden's, while the model uncertainty prices coincide with the solution to the malevolent player's decision process in the robust resource allocation problem. In the risk-return tradeoff under the approximating model but a preference for robustness, we must adjust the mean return  $\mu_r$  by  $\beta \cdot g^*$ .

Hansen and Jagannathan (1991) showed how the equity premium puzzle is revealed by the slope of the mean-standard deviation frontier. From the vantage-point of the approximating model, the instantaneous slope of the frontier is given by  $|\sigma_m + g^*|$ . One statement of the equity premium puzzle is that, in the absence of a preference for robustness, the factor risk prices implied by standard representative consumer models are small relative to the slope measured empirically. Robustness alters the slope prediction, because of a context specific form of pessimism that enters the model. A preference for robustness modifies the predicted mean-standard deviation tradeoff by adjusting factor prices to reflect a *market price of model uncertainty*  $|g_t^*|$ .

Thus a preference for robustness ends up attributing pessimistic beliefs to the decision maker. These pessimistic beliefs are then encoded in the asset prices. That pessimism

can help to resolve some asset return anomalies is not surprising. The contribution of our framework is that it restricts the form and magnitude of the pessimism. We will investigate how modest preferences for robustness can enhance theoretical values of market prices of risk. Again, we use statistical detection theory to give content to the term *modest*.

#### 2.4. Mathematical formulation

To study the implications of robustness, we will consider three families (semigroups) of operators indexed by elapsed calendar time. The first family consists of the conditional expectation operators associated with the Markov perfect equilibrium to the two-player game. The second family consists of pricing operators that assign prices to contingent claims in future time periods. The third family produces detection-error bounds for discriminating between competing models and is indexed by the sample interval used to classify models.

Each of these operator families has a generator that, in continuous time, governs the local evolution of the operators. As we will show, all three generators have the same mathematical structure. This connection links preferences for robustness to statistical detection and pricing.

Before our formal investigation of a continuous-time economy, we study robustness in decision-making using a recursive formulation in discrete time. We will then display simplifications that emerge when we move to a continuous-time diffusion formulation. As we will see, the recursive formulation of robustness coincides with a particular recursive formulation of risk preferences, an equivalence that has both conceptual and computational benefits. Moreover, this duality between risk aversion and robustness, as we model it, carries over to security market prices.

### 3. An operator and a probability bound

In this section, we study how to compute value functions under model misspecification in discrete time. We ignore control and study an exogenous state evolution equation. A Markov process governs the state vector. This process can be modeled by specifying a one-period conditional expectation operator:

$$\mathcal{T}(\phi)(y) = E[\phi(x_1) | x_0 = y].$$

The domain of the conditional expectation operator contains at least the space of bounded, continuous functions. It typically can be expanded to a larger space, but how depends on

the processes being modeled. Such extensions are relevant because we apply the conditional expectation operator to value functions that for some control problems can be unbounded.

Given a current-period reward function  $U(x)$  and a fixed Markov process, a value function  $W(x)$  for an infinite horizon solves the functional equation

$$W(x) = U(x) + \exp(-\delta) \mathcal{T}W(x) \quad (3.1)$$

where  $\delta$  is the subjective rate of discount and the period length is one.

Recursive utility formulations suggested by Epstein and Zin (1989) generalize (3.1) by replacing  $\mathcal{T}$  with an alternative transformation of the value function, which we denote  $\mathcal{R}$ . In the next section we aim to reinterpret the following specification of  $\mathcal{R}$  as incorporating a wish for robustness:

$$\mathcal{R}(W) = -\theta \log \left( \mathcal{T} \left[ \exp \left( \frac{-W}{\theta} \right) \right] \right). \quad (3.2)$$

As a formulation of recursive utility,  $\mathcal{R}$  was used by Weil (1993) in a closely related way. In the absence of discounting, replacing  $\mathcal{T}$  with  $\mathcal{R}$  in (3.1) delivers the risk sensitive evaluation used in control theory. The parameter  $\theta$  is restricted to be nonnegative and as it diverges to  $\infty$ ,  $\mathcal{R}$  reduces to the conditional expectation operator  $\mathcal{T}$ . The so-called risk sensitivity parameter is  $\frac{-1}{\theta}$ .

We want to interpret a recursion based on (3.2) in terms of a preference for robustness. We motivate this interpretation in terms of an exponential inequality that bounds the (conditional) tail probabilities of  $W$ . These tail probability bounds allow us to express a form of enhanced risk aversion generated by having the decision-maker care about more than just the conditional mean of the continuation value function. The tail probability bound comes from the theory of large deviation approximations. It uses the inequality:

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

as depicted in Figure 3.1. This inequality holds for any real number  $r$  and any  $\theta > 0$ . Let  $z$  denote the state tomorrow. Then computing expectations conditioned on the current state vector  $y$  yields:

$$\Pr\{W(z) \leq -r|y\} \leq E \left( \exp \left[ -\frac{W(z)}{\theta} \right] | y \right) \exp \left( -\frac{r}{\theta} \right),$$

---

<sup>14</sup> Weil (1993) used  $\mathcal{R}$  to make risk adjustments in a value function recursion that, in contrast to (3.1), is not additively separable. While there exists a transformation of the value function that has a recursion that is additively separable for Weil's formulation, the corresponding risk adjustment will be different.

<sup>15</sup> See Whittle (1990, 1996).


or

$$\log [\Pr\{W(z) \leq -r|y\}] \leq -\frac{1}{\theta} \mathcal{R}(W)(y) - \frac{r}{\theta}. \quad (3.3)$$

The first term on the right side of this inequality is independent of  $r$  but depends on  $\theta$ . We can express (3.3) as

$$\Pr\{W(z) \leq -r|y\} \leq \exp\{-\frac{1}{\theta} \mathcal{R}(W)(y)\} \exp\{-\frac{r}{\theta}\}. \quad (3.4)$$


Notice what this implies as  $r$  increases. Inequality (3.4) bounds the tail probability on the left by an exponential in  $r$ . It declines at rate  $\frac{-1}{\theta}$ ;  $\mathcal{R}$ , a function of  $\theta$ , influences the constant in the bound. Notice that decreasing  $\theta$  increases the exponential rate at which the bound sends the tail probabilities to zero. This expresses how a lower  $\theta$  heightens concern about tail events. This tells us how using  $\mathcal{R}$  to replace the mathematical expectation  $E$  in a typical Bellman equation enhances risk aversion.



**Figure 3.1:** Ingredients of large deviation bounds:  $\exp\left(\frac{-(W+r)}{\theta_i}\right)$  and  $\mathbf{1}_{\{W:W \leq -r\}}$  for  $\theta_1 = 1, r = 1$  and  $\theta_2 = 2, r = 1$ .

Figure 3.2 shows how  $\mathcal{R}$  induces additional caution about continuation utilities  $v$ . In the figure,  $E(v)$  is the expected utility of a gamble between two continuation utility levels  $v_2, v_1$  with  $v_2 > v_1$ . Where  $h(v)$  is a convex function, like  $\exp(-v/\theta)$  for  $0 < \theta < +\infty$ ,  $h^{-1}E(h(v)) < E(v)$ .

In the following, the operator  $\mathcal{R}$  will reappear but with a different interpretation, one reflecting a preference for robustness of continuation utility evaluations across a class of models. Notice that in the calculation of this section, there is only one model on the table, and  $\mathcal{R}$  is a device for injecting an additional adjustment to attitudes for risk



**Figure 3.2:** The function  $h^{-1}E(h(v))$  for  $h(v) = \exp(\frac{-v}{\theta})$ ,  $0 < \theta < +\infty$ .

above that captured by the curvature of  $U(x)$ . In the next section, we show that  $\mathcal{R}$  emerges as the outcome from a thought experiment in which the agent wants to evaluate continuation values under a neighborhood of models (transition measures). The parameter  $\theta$  will reappear as a measure of the size of the set of models.

#### 4. Misspecification in discrete time

We represent robustness in terms of a class of conditional expectation operators for perturbations of the original Markov process for the state. We start with an approximating model. It is a good approximation to what may be a more complicated model in feeding back on the history of the state. People agree on the approximating, perhaps from having observed common historical data. They worry about a common class of potential misspecifications in the form of perturbations around the approximating model. Agreement about the approximating model and the potential misspecifications is like imposing rational expectations. In this section, the approximating model is specified as a discrete-time Markov process. Later we study a continuous time setting. Instead of evaluating utility using only the approximating model, we use collection of candidate models that are centered around the approximating model.

#### 4.1. Markov perturbations

We start by introducing a particularly simple family of candidate models and discuss later what happens when we enrich this class. To enable us to use a formalization from large deviation theory for Markov processes (*e.g.* see Dupuis and Ellis, 1997), we generate candidate models as simple alterations of the approximating model's Markov transition probabilities. For a strictly positive function  $w$ , we form a distorted expectation operator

$$\mathcal{T}^w(\phi) = \frac{\mathcal{T}(w\phi)}{\mathcal{T}(w)}.$$

The operator  $\mathcal{T}^w$  gives rise to an alternative Markov specification for the state evolution. Given this construction,  $\frac{w(z)}{\mathcal{T}(w)(y)}$  is the Radon-Nikodym derivative of the altered or 'twisted' transition law with respect to the approximating model. Here  $z$  is a dummy variable used to index tomorrow's state and  $y$  is used to index today's state. Since  $\mathcal{T}^w$  is defined formally only for functions of tomorrow's state, it has the obvious extension to functions  $\psi(z, y)$  of both tomorrow's state and today's state. We let  $E^w(\cdot|y)$  denote this extension.

To embody the idea that the approximating model is good, we want to penalize discrepancies between a candidate alternative model and the approximating model. This can be done in a variety of ways. We measure discrepancy by 'relative entropy', defined as the expected value of the log-likelihood ratio (the log of the Radon-Nikodym derivative above), where expected value means the conditional expectation evaluated with respect to the density associated with the twisted or candidate alternative model (not the approximating model). For a candidate model indexed by  $w$ , this is

$$\begin{aligned} I(w)(y) &\equiv E^w \left[ \log \frac{w(z)}{\mathcal{T}(w)(y)} | y \right] \\ &= \mathcal{T}^w(\log w)(y) - \log[\mathcal{T}(w)(y)] \geq 0. \end{aligned}$$

Relative entropy is not a metric because it treats the approximating model and the candidate alternative model asymmetrically. This asymmetry emerges because the expectation is evaluated with respect to the 'twisted' or candidate distribution, not the one for the approximating model. Relative entropy is prominent in both information theory and large deviation theory and satisfies several attractive properties. (See Shore and Johnson, 1980 and Csiszar, 1991 for axiomatic justifications.) Not only is  $I(w)$  nonnegative, but

<sup>16</sup> We show later that preserving a Markov structure is not necessary.

<sup>17</sup> For readers of Dupuis and Ellis (1997, Chapter 1, Section 4), think of the transition density associated with  $\mathcal{T}$  as Dupuis and Ellis's  $\theta$ ; and think of  $\frac{w(z)}{\mathcal{T}(w)(y)}$  as Dupuis-Ellis's Radon-Nikodym derivative  $\frac{d\gamma}{d\theta}$ . For Dupuis and Ellis, relative entropy is  $\int \log \left( \frac{d\gamma}{d\theta} \right) d\gamma$ .

$I(w) = 0$  if  $w$  is constant. Substituting for  $\mathcal{T}^w$  gives:

$$\begin{aligned} I(w) &= \frac{\mathcal{T}[w \log(w)]}{\mathcal{T}(w)} - \log[\mathcal{T}(w)] \\ &= E \left[ \frac{w(z)}{E[w(z)|y]} \log \left( \frac{w(z)}{E[w(z)|y]} \right) \middle| y \right]. \end{aligned} \quad (4.1)$$

To summarize, we use relative entropy  $I(w)$  to measure how far candidate models (indexed by  $w$ ) are from the approximating model. In robust decision-making we pay particular attention to candidate models with small relative entropies.

#### 4.2. Value-function inequality for robustness

As an expression of a preference for robustness, we use the following recursion based on (3.1):

$$V(x) = U(x) + \exp(-\delta) \mathcal{R}W(x). \quad (4.2)$$

Here  $W(x)$  is a continuation value function and  $V(x)$  is a current value function. In this section, we describe an inequality that can give a bound on how much the conditional expectation of a value function will deteriorate across different probability specifications. Relative entropy and the penalty parameter  $\theta$  are the key objects in this bound.

Let  $W$  be a continuation value function, and  $\theta > 0$  a parameter. Consider the following problem:

PROBLEM A:

$$\inf_{w>0} J(w) \quad (4.3a)$$

where

$$J(w) \equiv \theta I(w) + \mathcal{T}^w(W). \quad (4.3b)$$

The first term on the right of (4.3b) is a weighted entropy measure and the second is the expectation of the value function using the twisted probability model indexed by  $w$ . The second term is the expectation of next period's value function when the current period's beliefs are as indexed by  $w$ . The objective is to find a worst-value model  $w$ , where the choice of departures  $w$  from the approximating model are penalized at a price of their relative entropy of  $\theta$ . Increasing the absolute magnitude of  $\theta$  increases the penalty for deviating from the approximating model.

---

<sup>18</sup> For Markov specifications with stochastically singular transitions,  $\frac{w(z)}{\mathcal{T}(w)(y)}$  may be one even when  $w$  is not constant. For these systems, we have in effect over parameterized the perturbations, although in a harmless way.

PROPOSITION 4.1. Assume that  $\mathcal{T}$  can be evaluated at  $\exp(-\frac{W}{\theta})$ . A solution to Problem A is:

$$w^* = \exp\left(-\frac{W}{\theta}\right),$$

which attains the minimized value

$$J(w^*) = \mathcal{R}(W)$$

where

$$\mathcal{R}(W) = -\theta \log\left(\mathcal{T}\left[\exp\left(-\frac{W}{\theta}\right)\right]\right)$$

The solution  $w^*$  is not unique because a scaled version of this function also obtains the same objective. However, the minimized value of the objective is unique and so is the associated probability law.

*Proof.* : To verify that  $w^*$  is the solution, write:

$$I(w) = I^*(w/w^*) + \frac{\mathcal{T}(w \log w^*)}{\mathcal{T}(w)} - \log \mathcal{T}(w^*)$$

where

$$I^*(w) = \frac{\mathcal{T}^*(w \log w)}{\mathcal{T}^*(w)} - \log \mathcal{T}^*(w)$$

and

$$\mathcal{T}^* \phi \equiv \frac{\mathcal{T}(w^* \phi)}{\mathcal{T}(w^*)}.$$

Notice that  $I^*$  is itself interpretable as a measure of relative entropy and hence  $I^*(w/w^*) \geq 0$ . Thus the criterion  $J$  satisfies the inequality:

$$\begin{aligned} J(w) &= \theta \left[ I^*(w/w^*) + \frac{\mathcal{T}(w \log w^*)}{\mathcal{T}(w)} - \log \mathcal{T}(w^*) \right] + \mathcal{T}^w(W) \\ &\geq \theta \left[ \frac{\mathcal{T}(w \log w^*)}{\mathcal{T}(w)} - \log \mathcal{T}(w^*) \right] + \mathcal{T}^w(W) \\ &= -\theta \log \mathcal{T} \left[ \exp\left(-\frac{W}{\theta}\right) \right] \\ &= J(w^*). \end{aligned}$$

■

When we discussed equations (3.2) and (3.3), we interpreted  $\mathcal{R}$  in terms of making an extra risk adjustment, and we described an associated bound on the tail probabilities

---

<sup>19</sup> This proof mimics the proof of Proposition 1.4.2 in Dupuis and Ellis (1997), but is included for completeness.

of the continuation value function  $W$ . Equation (4.3b) implies an inequality in terms of robust evaluations of value functions.

COROLLARY 4.1. *The conditional expectation of the value function  $W$  evaluated under the transition  $\mathcal{T}^w$  satisfies the bound*

$$\mathcal{T}^w(W) \geq \mathcal{R}(W) - \theta I(w). \quad (4.4)$$

*Proof.* This follows immediately from  $J(w^*) = \mathcal{R}(W)$  and the definition of  $J(w)$ . ■

The first term on the right depends on  $\theta$ , but not on the alternative model parameterized by  $w$ . The second term is  $-\theta$  times entropy, our measure of discrepancy between models. Thus, inequality (4.4) identifies  $\theta$  as a type of *utility price of robustness*. The robust value function  $W$  solves the functional equation:

$$\begin{aligned} W(x) &= \inf_w (U(x) + \exp(-\delta) [\theta I(w) + \mathcal{T}^w(W)(x)]) \\ &= U(x) + \exp(-\delta) \mathcal{R}(W)(x), \end{aligned}$$

which displays the robust interpretation of adjusting for the riskiness of the future value function.

### 4.3. Richer class of perturbations

While formula (4.1) exploits the specific Markov structure of the candidate model, the relative entropy measure has a straightforward extension to a much richer class of perturbations. For instance, suppose that  $w$  is a strictly positive function that is allowed to depend on both  $z$  and  $y$ . In this case we extend the measure of entropy to:

$$I(w) = E \left[ \frac{w(z, y)}{E[w(z, y) | y]} \log \left( \frac{w(z, y)}{E[w(z, y) | y]} \right) \right].$$

The solution to the extended version of Problem A remains the same. In other words, it suffices for the ‘control’  $w$  to depend only on the state tomorrow. Notice the perturbation just described preserves the first-order Markov structure. That is, the ‘twisted’ process remains a first-order Markov process.

First-order Markov perturbations are special. So it is of interest to extend the class of perturbations further. We construct a bigger class of perturbations by starting with positive (and appropriately measurable) functions that depend on the entire past history of the Markov state for the approximating model. Let  $t + 1$  denote tomorrow’s date. Form the Radon-Nikodym derivative as:

$$h_{t+1} = \frac{w(z_{t+1}, z_t, z_{t-1}, \dots)}{E[w(z_{t+1}, z_t, z_{t-1}, \dots) | z_t, z_{t-1}, \dots]}.$$

In other words, the time  $t + 1$  Radon-Nikodym derivative is a strictly positive random variable that is measurable with respect to the current and past values of the Markov state (for the approximating model). This random variable is constrained to have mean one conditioned on  $z_t, z_{t-1}, \dots$ . The time  $t$  conditional entropy is then:

$$I_t(h_{t+1}) = E[h_{t+1} \log h_{t+1} | z_t, z_{t-1}, \dots].$$

The time  $t$  counterpart to the objective function  $J$  for Problem A is:

$$J_t(h_{t+1}) = \theta I_t(h_{t+1}) + E[h_{t+1} W(z_{t+1}) | z_t, z_{t-1}, \dots].$$

Given that the process  $\{z_t\}$  is Markov under the approximating model, the solution to the counterpart to control problem A is to let the ‘control’  $h_{t+1}$  be a function of the Markov state  $z_{t+1}$  alone. Thus our solution to Problem A extends to this richer class of perturbations, including ones for which the time  $t+1$  distortion fails to preserve the Markov structure. The essential feature of these perturbations is that the transition probability distribution for the candidate model be absolutely continuous with respect to the transition probabilities of the approximating model; this manifests itself in the restriction that  $h_{t+1} > 0$ . Absolute continuity keeps the log likelihood ratio well defined.

## 5. Misspecification in continuous time

We now study the continuous-time counterpart to Problem A. A continuous-time formulation simplifies the analysis. For example, we will show that when the approximating model is a diffusion, a specification of  $w$  just introduces a nonzero drift specification into Brownian motion increments that drive the model. Some Markov jump processes and mixed jump diffusion models can also be handled. In the case of a Markov jump process, the misspecification  $w$  alters both the jump intensities and the chain probabilities that dictate jump locations when a jump takes place.

The operator formulation of continuous-time Markov processes entails specifying an infinitesimal generator  $\mathcal{A}$ . The conditional expectations are then given heuristically by:

$$\mathcal{T}_s = \exp(s\mathcal{A}).$$

This construction is formalized using a Yosida approximation. The domains of the conditional expectation operators  $\mathcal{T}_s$  will contain at least the space  $\hat{C}$  of continuous functions

$\phi$  with limit zero as  $|x| \rightarrow \infty$ ; these are equipped with the sup norm. The family or semigroup of conditional expectation operators  $\{\mathcal{T}_s \geq 0\}$  is assumed to be a Feller semigroup, which among other things implies that

$$\lim_{s \downarrow 0} \mathcal{T}_s = \mathcal{I}$$

where  $\mathcal{I}$  denotes the identity operator. (See Ethier and Kurtz, 1986, Chapter 1, for a general discussion of semigroups and Chapter 4 for a discussion of Feller semigroups.)

### 5.1. Examples

When the approximating model is a diffusion, the generator is a second-order differential operator:

$$\mathcal{A}\phi = \mu \cdot \frac{\partial \phi}{\partial x} + \frac{1}{2} \text{trace} \left( \Sigma \frac{\partial^2 \phi}{\partial x \partial x'} \right),$$

where the coefficient vector  $\mu$  is the drift of the diffusion and the coefficient matrix  $\Sigma$  is the diffusion coefficient. The corresponding stochastic differential equation is:

$$dx_t = \mu(x_t) dt + \Lambda(x_t) dB_t$$

where  $\{B_t\}$  is a multivariate standard Brownian motion and  $\Lambda\Lambda' = \Sigma$ . In this case the generator is not a bounded operator and its domain excludes some of the functions in  $\hat{C}$ . Nevertheless, the generator's domain includes at least functions that are twice differentiable and have compact support. We have reason to extend this domain in our later analysis.

When the approximating model is a Markov jump process, the generator can be represented as:

$$\mathcal{A}\phi = \lambda [\mathcal{S}\phi - \phi]$$

where the coefficient  $\lambda$  is a Poisson intensity parameter that sets the jump probabilities and  $\mathcal{S}$  is a conditional expectation operator that encodes the transition probabilities conditioned on a jump taking place. The intensity parameter can depend on the Markov state and is bounded and nonnegative. In this case the domain of both the generator and the semigroup can be extended to the space of bounded Borel measurable functions, again equipped with the sup norm. The generator is a bounded operator on this enlarged space, and the exponential formula for recovering the conditional expectation operators can be defined formally using the standard series expansion of an exponential function. (See Ethier and Kurtz, 1986, Chapter 4).

## 5.2. Markov Process Perturbations

Next we define a family of perturbations of the approximating Markov process. We start with the perturbations used in discrete time and taking limits. In so doing, we will construct ‘perturbed’ generators. At the outset we proceed heuristically, partly imitating our discrete-time investigation. Let  $w$  be a positive function. For a small time interval  $\epsilon$  form:

$$\mathcal{T}_\epsilon^w \phi = \frac{\mathcal{T}_\epsilon(w\phi)}{\mathcal{T}_\epsilon(w)}.$$

To construct a ‘twisted’ generator of continuous-time process associated with  $w$ , we compute:

$$\begin{aligned} \mathcal{A}^w \phi &\equiv \lim_{\epsilon \downarrow 0} \frac{\mathcal{T}_\epsilon^w \phi - \phi}{\epsilon} \\ &= \frac{\mathcal{A}(w\phi) - \phi \mathcal{A}(w)}{w}. \end{aligned}$$

The resulting formula is well defined provided that  $w$  and  $w\phi$  are both in the domain of the generator. Since our candidate generator is obtained as a limit of discrete-time generators, we must verify that the  $\mathcal{A}^w$  actually generates a Feller semigroup. One way to do this is to verify the postulates of the Hille-Yosida Theorem (Ethier and Kurtz, 1986, page 165). Instead, in later sections we will explore this construction directly for the approximating model processes that interest us.

To study entropy penalties in continuous time, let  $I_\epsilon(w)$  be the small increment counterpart to the discrete-time entropy measure  $I(w)$ :

$$I_\epsilon(w) = \frac{\mathcal{T}_\epsilon[w \log(w)]}{\mathcal{T}_\epsilon(w)} - \log[\mathcal{T}_\epsilon(w)].$$

Although  $I_\epsilon(w)$  converges to zero as  $\epsilon$  declines to zero, its ‘derivative’ is nondegenerate:

$$I'(w) \equiv \lim_{\epsilon \downarrow 0} \frac{I_\epsilon(w)}{\epsilon} = \frac{\mathcal{A}[w \log(w)]}{w} - \frac{\log w}{w} \mathcal{A}(w) - \frac{1}{w} \mathcal{A}(w).$$

Since  $I_\epsilon(w)$  is nonnegative, so is  $I'(w) \geq 0$ , and  $I'(w)$  equals zero when  $w$  is constant.

Combining these two limiting operators, we can construct a continuous-time counterpart to criterion  $J$ :

$$J'(w) = \theta I'(w) + \mathcal{A}^w W. \tag{5.1}$$

Thus to perform a robust version of continuous-time value function updating we are led to solve:

PROBLEM B

$$\inf_{w>0, w \in D} J'(w), \tag{5.2a}$$

where  $D$  is constructed so that both  $w$  and  $wW$  are in the domain of the generator.

PROPOSITION 5.1. *The minimizer of Problem B is*

$$w^* = \exp\left(-\frac{W}{\theta}\right). \quad (5.2b)$$

*The optimized value of the criterion is:*

$$J'(w^*) = -\theta \frac{\mathcal{A}\left[\exp\left(-\frac{W}{\theta}\right)\right]}{\exp\left(-\frac{W}{\theta}\right)}. \quad (5.2b)$$

*Proof.* To verify the conjectured solution, we imitate the argument from discrete time. We change the reference point for the relative entropy measure by using the ‘twisted’ generator  $\mathcal{A}^* = \mathcal{A}^w$  with  $w = w^*$  in place of  $\mathcal{A}$ :

$$I^{*'}(w) \equiv \frac{\mathcal{A}^*[w \log(w)]}{w} - \frac{\log w}{w} \mathcal{A}^*(w) - \frac{1}{w} \mathcal{A}^*(w).$$

It can be verified that

$$\begin{aligned} I^{*'}(w/w^*) &= I'(w) - I'(w^*) + \frac{1}{\theta} \mathcal{A}^w(W) - \frac{1}{\theta} \mathcal{A}^*(V) \\ &= \theta^{-1} [J'(w) - J'(w^*)]. \end{aligned}$$

Since the left-hand side is always nonnegative and  $\theta$  is positive, it follows that  $w^*$  solves Problem B. ■

COROLLARY 5.1. *The infinitesimal generator for the value function under model  $w$  obeys*

$$\mathcal{A}^w(W) \geq -\theta \frac{\mathcal{A}\left(\exp\left(\frac{-W}{\theta}\right)\right)}{\exp\left(\frac{-W}{\theta}\right)} - \theta I'(w). \quad (5.3)$$

*Proof.* This follows from the definition of (5.1) and result (5.2b). ■

This is the analogue of (4.4). The first term depends on  $\theta$  but is independent of the perturbation  $w$ . Here  $\theta$  is the price in terms of the assured infinitesimal generator for value of increasing entropy at rate  $I'(w)$ .

---

<sup>20</sup> Given the presence of the inf instead of a min, the restrictions on  $w$  are not problematic. However, restricting  $V$  to be in the domain of the generator is sometimes too severe. This restriction can be relaxed by instead using the extended generator.

### 5.3. A Robust Value Function Recursion

We now revisit the question we posed at the outset of section 2. In discrete time, the usual fixed point problem for evaluating a time-invariant control law under discounting and an infinite horizon is:

$$V(x) = \epsilon U(x) + \exp(-\epsilon\delta) \mathcal{T}_\epsilon V(x)$$

where  $U$  is the current-period reward function and  $\delta$  is the subjective rate of time discount. We obtain a continuous-time counterpart by subtracting  $V$  from both sides, dividing by  $\epsilon$ , and taking limits. This results in:

$$0 = U - \delta V + \mathcal{A}V.$$

When we introduce adjustments for model misspecification, we modify this to be

$$0 = U - \delta V - \theta \frac{\mathcal{A} \left[ \exp\left(-\frac{V}{\theta}\right) \right]}{\exp\left(-\frac{V}{\theta}\right)},$$

which can be viewed as one of the utility recursions studied by Duffie and Epstein (1992).

The connection between risk sensitivity and robustness is also used in a literature on risk-sensitive control (*e.g.* see James (1992) and Runolfsson (1994)). In contrast to that work, we consider control problems with discounting, which leads to a recursive formulation of risk sensitivity. The resulting recursion is the continuous-time generalization of one studied by Hansen and Sargent (1995).

## 6. Continuous-time diffusions

In this section we study the solution to Problem B when the underlying approximating model is a diffusion. We then describe a larger class of perturbations in continuous-time that are consistent with our entropy-based notion of robustness.

### 6.1. Markov Perturbations

We start by specializing the class of perturbations to the diffusion model. Hence, we take the generator to be:

$$\mathcal{A}\phi = \mu \cdot \frac{\partial\phi}{\partial x} + \frac{1}{2}\text{trace} \left( \Sigma \frac{\partial^2\phi}{\partial x \partial x'} \right),$$

and, as before, we consider a twisted generator:

$$\mathcal{A}^w\phi = \frac{\mathcal{A}(w\phi) - \phi\mathcal{A}(w)}{w}.$$

Note that by the product rule for first and second derivatives:

$$\mathcal{A}(w\phi) = w\mathcal{A}\phi + \phi\mathcal{A}w + \left( \frac{\partial w}{\partial x} \right)' \Sigma \left( \frac{\partial\phi}{\partial x} \right).$$

Thus the twisted generator can be represented by

$$\mathcal{A}^w\phi = \mathcal{A}\phi + g'\Lambda' \left( \frac{\partial\phi}{\partial x} \right)$$

where we have defined  $g$  to be the logarithmic derivative

$$g \equiv \Lambda' \frac{\partial \log w}{\partial x}.$$

Therefore this form of model distortion is equivalent to adding a state dependent drift  $g$  to the Brownian motion of the approximating model:

$$dx_t = \mu(x_t) dt + \Lambda(x_t) [g(x_t) dt + dB_t]. \quad (6.1)$$

The corresponding continuous-time measure of relative entropy is:

$$\begin{aligned} I'(w) &= \frac{\mathcal{A}[w \log(w)]}{w} - \frac{\log w}{w} \mathcal{A}(w) - \frac{1}{w} \mathcal{A}(w) \\ &= \frac{1}{2} g'g. \end{aligned}$$

Using our representations for the twisted Markov model and of relative entropy, we can rewrite Problem B as:

PROBLEM C:

$$\min_g \left[ \frac{\theta}{2} g'g + \mathcal{A}V + g'\Lambda' \left( \frac{\partial V}{\partial x} \right) \right].$$

This problem has solution:

$$g^* = -\frac{1}{\theta} \Lambda' \left( \frac{\partial V}{\partial x} \right).$$

The robust version of the value function recursion is now:

$$U - \delta V + \mathcal{A}V - \frac{1}{2\theta} \left( \frac{\partial V}{\partial x} \right)' \Sigma \left( \frac{\partial V}{\partial x} \right) = 0. \quad (6.2)$$

Runolfsson (1994) deduced the  $\delta = 0$  (ergodic control) counterpart to (6.2) to obtain a robust interpretation of risk sensitivity. Partial differential equation (6.2) is also a special case of the equation system that Duffie and Epstein (1992) and Duffie and Lions (1992) analyze for stochastic differential utility. They showed that for diffusion models, the recursive utility generalization introduces a variance multiplier, which can be state dependent. The counterpart to this multiplier in our setup is state independent and equal to the so called risk sensitivity parameter  $\frac{-1}{\theta}$ . Notice that variance multiplier is used to weight the quadratic form  $\left( \frac{\partial V}{\partial x} \right)' \Sigma \left( \frac{\partial V}{\partial x} \right)$ , which is the local variance of the value function process  $\{V(x_t)\}$ . The risk interpretation of (6.2) is that the decision maker is concerned not only about the local mean of the continuation value process, but also about the local variance of this process.

Though we are interested in this variance multiplier for its role in restraining entropy, its mathematical connection to risk sensitivity and recursive utility lets us draw on a set of analytical results from those literatures.

### 6.2. Enlarging the class of perturbations

So far we have looked only at distortions that append a drift  $\{\int_0^t g(x_s)ds\}$  to the Brownian motion, which preserves the Markov specification. Alternatively, we could follow James (1992) by considering path dependent specifications of the drift of the Brownian motion  $\int_0^t g_s ds$ , where  $g_s$  is constructed as a general function of past  $x$ 's. Thus the state evolution is

$$dx_t = \mu(x_t) dt + \Lambda(x_t) [g_t dt + dB_t]. \quad (6.3)$$

The 'control' vector process  $\{g_t\}$  is suitably adapted to the filtration generated by the underlying Brownian motion vector. The altered process may no longer be Markovian.

To incorporate this form of robustness, we alter the time  $t$  reward function of the decision-maker by augmenting  $U(x_t)$  with  $\frac{\theta}{2}|g_t|^2$ . The term  $(1/2)g_t'g_t$  is a time  $t$  measure of relative entropy but for a more general family of perturbations. The resulting objective of the robust decision-maker has the representation

$$E \int_0^\infty \exp(-\delta t) \left[ U(x_t) + \frac{\theta}{2} g_t' g_t \right] dt. \quad (6.4)$$

Given the Markov structure of this control problem, its solution can be represented as a time-invariant function of the state vector  $x_t$ , which we denote  $g_t^* = g^*(x_t)$ . Notice

that (6.1) and (6.3) coincide under a time invariant control law. The resulting Bellman equation is identical to (6.2), and hence the optimal control satisfies

$$g_t^* = -\frac{1}{\theta} \Lambda(x_t)' \left( \frac{\partial V}{\partial x} \right) (x_t).$$

This form of perturbation looks very special. Our discussion of discrete-time models might have led us to expect a large class of perturbations with finite (relative entropy). However, Girsanov's Theorem tells us that there is no scope for enlargement: absolute continuity is a potent restriction for continuous-time diffusions. Thus, by looking at diffusions we simplify the dichotomy between models that are statistically close and those that are not. In order for models to be close in the sense of relative entropy, their diffusion matrices must coincide. This is because instantaneous covariance matrices can be inferred with arbitrary accuracy from high frequency data. Estimating the drift is more difficult. Hence we focus our concerns about model misspecification on the drift component. These simple drift distortions produce a rich collection of candidate models.

### 6.3. Adding controls to the original state equation

We now add a control vector to the original state evolution equation

$$dx_t = \mu(x_t, i_t) dt + \Lambda(x_t, i_t) (g_t dt + dB_t) \quad (6.5)$$

and to the return function  $U(x, i)$  where  $i_t$  is a control vector. Consider a time-invariant control of the form  $i_t = f(x_t)$ . The value function for the robust control problem satisfies (6.2), except that now the differential operator  $\mathcal{A}$  and the reward function  $U$  depend on the control law  $f$ :

$$U(\cdot, f) - \delta V + \mathcal{A}(f) V - \frac{1}{2\theta} \left( \frac{\partial V}{\partial x} \right)' \Sigma(\cdot, f) \left( \frac{\partial V}{\partial x} \right) = 0.$$

As we have seen, for a fixed control law  $f$ , this equation emerges from a single agent robust control problem. Thus we can solve for the optimal (or robust) control law by solving a two-player Markov game:

$$\max_f \min_g \left[ U(x, f) - \delta V(x) + \mathcal{A}(f) V(x) + \left( \frac{\partial V}{\partial x}(x) \right)' \Sigma(x, f) u + \frac{\theta}{2} g' g \right] = 0. \quad (6.6)$$

---

<sup>21</sup> See, *e.g.*, Karatzas and Shreve (1991, pages 190-196), Durrett, (1996, pages 90-93) and Dupuis and Ellis (1997, pages 123-125).

#### 6.4. Two linear technologies

We now consider a special case of the stochastic growth model described in the overview. There are two linear technologies. The first technology is risk-free with return  $r$ . The second technology is risky with an expected return of  $r + \lambda$  where  $\lambda > 0$ . The instantaneous variance of the second technology is  $\sigma^2$ . We assume that  $r, \lambda$  and  $\sigma$  are time-invariant. Let  $\psi$  denote the fraction of aggregate capital invested in the risky technology. Aggregate capital evolves as

$$dk_t = \psi k_t dy_t + (rk_t - c_t) dt$$

where

$$dy_t = \lambda dt + \sigma (g_t dt + dB_t) \quad (6.7)$$

and where  $B_t$  is a scalar Brownian motion. Here  $dy_t$  gives the excess return of the risky technology. This economy is a special case of example two from section 2 with the  $\alpha_i$ 's equal to one and the  $\delta_i$ 's equal to zero.

The Bellman equation for the planner is given by equation (6.6). The state is aggregate capital, the control vector is  $[c \ \psi]'$ , and

$$\begin{aligned} \Sigma &= \sigma^2 \psi^2 k^2, \\ \mathcal{A}V &= \frac{\partial V}{\partial k} (\psi k \lambda + rk - c) + \frac{1}{2} \sigma^2 k^2 \frac{\partial^2 V}{\partial k^2}. \end{aligned}$$

Notice that  $\Sigma$  depends on aggregate capital and the control vector. The first order conditions are

$$U'(c) = \frac{\partial V}{\partial k} \quad (6.8)$$

$$\psi = -\frac{1}{k} \left[ \frac{\frac{\partial V}{\partial k}}{\frac{\partial^2 V}{\partial k^2} - \frac{1}{\theta} \left(\frac{\partial V}{\partial k}\right)^2} \right] \frac{\lambda}{\sigma^2}. \quad (6.9)$$

When  $-\frac{1}{\theta} = 0$  in equation (6.9), we obtain the usual result that the fraction of capital allocated to the risky technology is the inverse of the Arrow-Pratt measure of relative risk-aversion times  $\sigma^{-2}\lambda$ . Agents with a preference for robustness act as though they have an Arrow-Pratt risk-aversion of

$$-k \frac{\frac{\partial^2 V}{\partial k^2}}{\frac{\partial V}{\partial k}} + \frac{k}{\theta} \frac{\partial V}{\partial k}. \quad (6.10)$$

This is the usual formula for the Arrow-Pratt measure of risk-aversion plus  $\frac{k}{\theta} \frac{\partial V}{\partial k}$ . It can be shown for a discrete time economy with value function  $V$  that (6.10) is the appropriate generalization of the inverse of the Arrow-Pratt measure of relative risk-aversion for risk-sensitive preferences. When agents want robustness, we will call equation (6.10) *effective risk-aversion*.

**Evan: the figures need fixing: where are the right and left panels you refer to ?**

We now display some example solutions to the optimal resource allocation problem. In all the examples, we assume that a representative agent has a reward function of the form

$$U(c) = \frac{c^{1-\gamma}}{1-\gamma}.$$

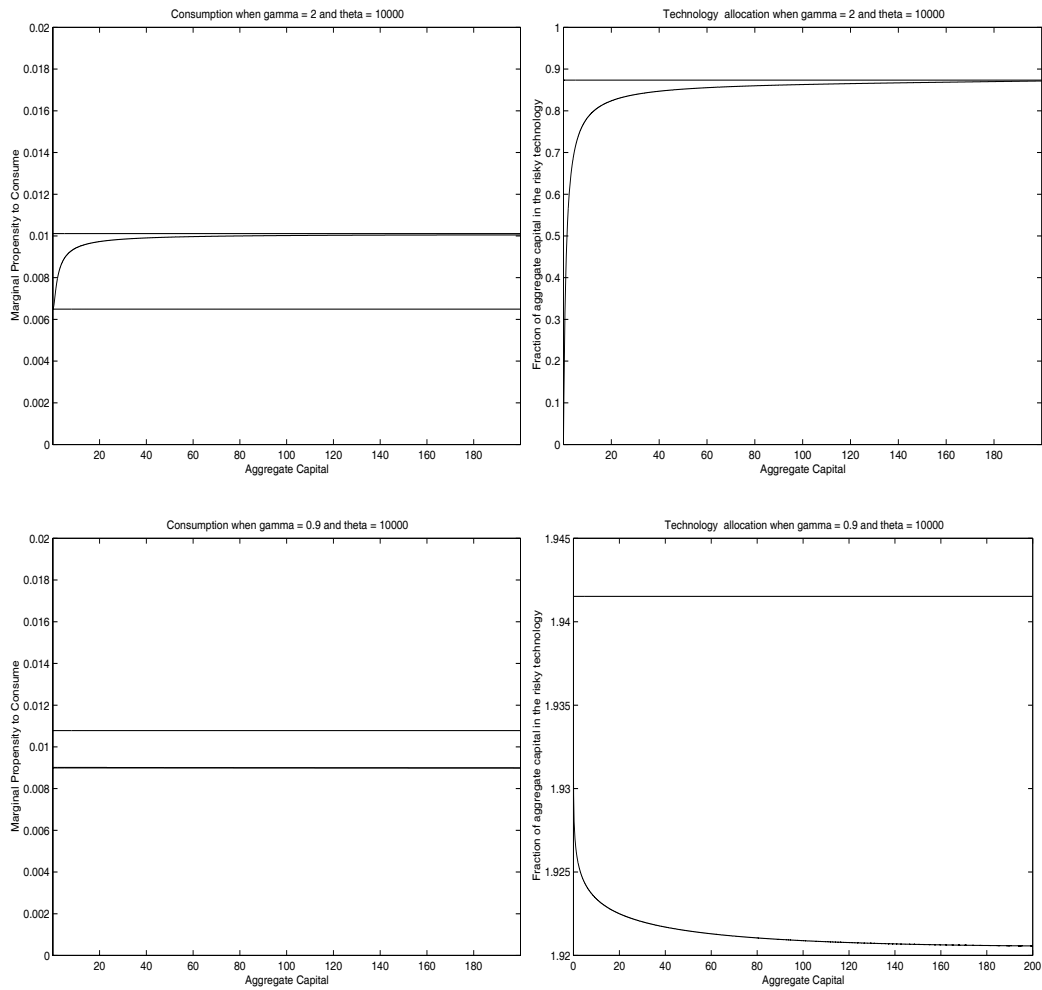
Let

$$r = .002975, \quad \lambda = .0166, \quad \sigma^2 = .0975, \quad \delta = 0.01, \quad \theta = 10000.$$

The values for  $r$ ,  $\lambda$  and  $\sigma$  were designed to roughly be consistent with quarterly U.S. data. The value of  $\delta$  was chosen arbitrarily. When  $\gamma \neq 1$  and  $\theta < \infty$ , the optimal decision rules are nonlinear and there are no known analytical formulas for the rules.

The right graphs in Fig. 6.1 display the allocation rules for two values of  $\gamma$ . The fraction of aggregate capital invested in the risky technology is bounded above by the fraction invested in the risky technology when the agents have not preferences for robustness. The fraction is increasing in aggregate capital when  $\gamma > 1$  and decreasing in aggregate capital when  $\gamma < 1$ . When  $\gamma > 1$ , as the aggregate capital increases, the fraction converges to the fraction invested in the risky technology when the agents are not robust. From formula (6.9), the left graphs in Fig. 6.1 can be interpreted as the inverse of the planner's effective relative risk-aversion times  $\sigma^{-2}\lambda$ . Thus we see that effective risk-aversion is increasing in aggregate capital when  $\gamma > 1$  and decreasing in aggregate capital when  $\gamma < 1$ .

The left graphs in Fig. 6.1 displays the consumption choice rules for two values of  $\gamma$ . The consumption choice rules are bounded above by the rules in an economy with no preference for robustness. Consumption divided by aggregate capital is increasing in aggregate capital when  $\gamma > 1$  but it is nearly globally linear and converges quickly to the level of consumption in an economy without robustness. When  $\gamma < 1$ , consumption is almost indistinguishable from a linear function. In this case, it appears that the consumption function is globally almost collinear with the consumption function with no preference for robustness, where  $\lambda = 0$ .



**Figure 6.1: Consumption and allocation decision rules.** The first row displays the consumption and allocation decision rules when  $\gamma = 2$ . The second row displays the consumption and allocation decision rules when  $\gamma = 0.9$ . The lower straight lines on the left graphs are the marginal propensity to consume when there is no robustness and  $\lambda = 0$ . (Note that in the bottom left graph, the lower dashed line is almost perfectly collinear with the optimal robust decisions.) The upper straight lines on the left graphs are the marginal propensity to consume when there is no robustness. The straight line in the right graphs is the allocation choice when there is

## 7. Equilibrium pricing

Under a preference for robustness, Hansen, Sargent, and Tallarini (1999) obtained an approximate additive decomposition of the market price of risk for a linear-quadratic discrete time model. By using a continuous time formulation, this section derives an exact counterpart to their representation that dispenses with the linear-quadratic assumptions.

When production and capital accumulation are feasible, activating a preference for robustness alters a planner's optimal allocation. Equilibrium prices emerge from the robust planner's attitudes about marginal transformations of consumption across dates and states, evaluated at the optimal robust allocation. This observation allows us to compute equilibrium asset prices under a preference for robustness by treating the allocation as though it were exogenous or had emerged from a pure endowment economy. This modeling device allows us to deduce simple formulas for model uncertainty premia in security market prices.

We write the consumption evolution as

$$\begin{aligned} dx_t &= \mu^*(x_t) dt + \Lambda^*(x_t) dB_t \\ c_t &= \psi(x_t) \end{aligned}$$

where  $c_t$  is the solution for log consumption from the robust resource allocation problem,  $\mu^*$  is the drift coefficient when the optimal robust control law is imposed, and

$$\Sigma^* = \Lambda^* \Lambda^{*'}$$

is the corresponding diffusion matrix.

Formally, we proceed in three steps. First we form operators that assign current period prices to claims on consumption in future periods. Next we show how these prices would be related to the intertemporal marginal rates of substitution if we had not posited a preference for robustness, an exercise that reproduces results from Breeden's (1979) analysis of the consumption-based asset pricing model. Finally, we show how prices are altered when we introduce a preference for robustness. This final step allows us to formulate a pricing premium due to model uncertainty.

---

<sup>22</sup> When production and capital accumulation are introduced, robustness will also alter the intertemporal consumption allocation through a precautionary impact on investment behavior. Our price decomposition, but not our equilibrium price computation, abstracts from the change in allocations due to robustness.

### 7.1. Representing Pricing Operators

As with conditional expectation operators, equilibrium security prices in Markov models can be represented using the mathematical machinery of positive operators. For the case of Markov diffusions, we construct a family (semigroup) of operators that assigns date  $t$  prices to contingent claims at date  $t+s$ . The domains of these operators are appropriately restricted functions of the Markov state. We consider contingent consumption claims at date  $t+s$  that can be expressed as functions  $\phi$  of the state  $x_{t+s}$ . The pricing operator (for interval  $s$ ) assigns a date  $t$  price to this claim.

Since the solution to the robust resource allocation problem is a Markov diffusion process, the family of pricing operators can be represented using stochastic discount factors via:

$$\mathcal{P}_s(\phi)(y) = E[\phi(x_{t+s}) q_{t,t+s} | x_t = y] \quad (7.1)$$

where the stochastic discount factor between dates  $t$  and  $t+s$  is the positive random variable:

$$q_{t,t+s} = \exp\left(-\int_t^{t+s} [\eta(x_u) du + \lambda(x_u) \cdot dB_u]\right). \quad (7.2)$$

It is convenient to decompose the stochastic discount factor into two multiplicative components:

$$q_{t,t+s} = q_{t,t+s}^1 q_{t,t+s}^2$$

where

$$\begin{aligned} q_{t,t+s}^1 &\equiv \exp\left(-\int_t^{t+s} \left[\eta(x_u) - \frac{1}{2}\lambda(x_u) \cdot \lambda(x_u)\right] du\right) \\ q_{t,t+s}^2 &\equiv \exp\left(-\int_t^{t+s} \left[\frac{1}{2}\lambda(x_u) \cdot \lambda(x_u) du + \lambda(x_u) \cdot dB_u\right]\right). \end{aligned}$$

As we will see, each component of the stochastic discount factor contributes to pricing in an interpretable way.

This stochastic discount factor representation (7.1) is known to be valid when the following two restrictions are met:

ASSUMPTION 7.1.  $\eta - \frac{1}{2}\lambda'\lambda$  is bounded from below.

ASSUMPTION 7.2.  $E(\exp[\int_t^{t+s} \lambda(x_u) \cdot \lambda(x_u) du] | x_t = y) < \infty$  for all  $y$  and  $s$ .

The first assumption is a technical condition imposed to make the additional  $du$  discounting using  $q_{t,t+s}^1$  well defined for a large collection of contingent claims (functions  $\phi$ .) As we will show, this first component gives rise to the instantaneous risk-free interest rate and hence to the  $dt$  prices.

The second assumption is the Novikov (1972) condition used to guarantee that  $\{q_{t,t+s}^2 : s \geq 0\}$  is an exponential martingale. The  $du$  correction is the familiar lognormal adjustment needed to compensate for exponentiation. Use of  $q_{t,t+s}^2$  alone gives rise to the familiar risk-neutral transformation of the probabilities for pricing derivative claims. Since  $q_{t,t+s}^2$  is positive and has conditional mean equal to unity, the pricing operator constructed with only this component of the discount factor is also a conditional expectation operator, but with a transformed (risk neutral) probability distribution. As we will show, this component gives rise to the so-called factor risk prices, in effect the prices of the Brownian motion increments.

## 7.2. Incremental Prices

We use the generator  $\mathcal{B}$  for the family (semigroup) of pricing operators to define the local prices. We then consider how this generator is related to the generator  $\mathcal{A}$  for the semigroup of conditional expectation operators. We will study the difference between the two generators to provide formulas for the prices of the Brownian motion increments and the prices for the  $dt$  increments.

The generator  $\mathcal{B}$  of the pricing semigroup is defined formally as the limit

$$\mathcal{B}\phi = \lim_{s \downarrow 0} \frac{\mathcal{P}_s\phi - \phi}{s}$$

for an (appropriately restricted) functions  $\phi$ . This operator generates the semigroup via the exponential formula:

$$\mathcal{P}_s \approx \exp(s\mathcal{B}),$$

which can be justified rigorously using a Yosida approximation.

Given the diffusion formulation of the stochastic discount factor, it can be shown that

$$\mathcal{B}\phi(y) = \mathcal{A}\phi - \left(\eta - \frac{1}{2}\lambda \cdot \lambda\right)\phi - \frac{\partial\phi'}{\partial x}\Lambda^*\lambda. \quad (7.3)$$

Thus generator  $\mathcal{B}$  of the pricing operators includes the two additional terms:

$$- \left(\eta - \frac{1}{2}\lambda \cdot \lambda\right)\phi - \frac{\partial\phi'}{\partial x}\Lambda^*\lambda \quad (7.4)$$

*vis a vis* the generator of the family of conditional expectation operators.

To interpret the first term in (7.4), we compute the instantaneous rate of interest on a riskless bond. We calculate this rate by evaluating the limit of the logarithms of the

discount bond prices:

$$\begin{aligned}\rho &\equiv -\lim_{s \downarrow 0} \frac{\log(\mathcal{P}_s 1)}{s} \\ &= -\mathcal{B}1 \\ &= \eta - \frac{1}{2} \lambda \cdot \lambda.\end{aligned}$$

Thus the coefficient multiplying the level term  $\phi$  for the pricing generator is the negative of the instantaneous risk-free rate. This level term is included because, unlike a conditional expectation operator, the pricing operator must discount constant functions (riskless payoffs). In effect,  $\rho$  prices the  $dt$  increments.

The coefficient  $\Lambda^* \lambda$  multiplying the first derivative  $\frac{\partial \phi}{\partial x}$  gives the change in the drift term when moving from the conditional expectation operator to the pricing operator. This transformation is the familiar adjustment from the pricing of derivative claims via the use of a risk neutral probability distribution. The vector  $\lambda$  is commonly referred to as the vector of *factor risk prices*. Note that  $\frac{\partial \phi'}{\partial x} \Lambda^* dB$  is the Brownian increment associated with the payoff process  $\{\phi(x_t)\}$ . Thus  $\frac{\partial \phi'}{\partial x} \Lambda^* \lambda$  is the dot product between the factor weight for the Brownian increments of the payoff and the *price vector*  $\lambda$ .

In summary, the incremental prices are given by the instantaneous mean of the payoff adjusted by the instantaneous risk-free rate and by a price correction for the factor loadings on the Brownian increments. Prices over any finite horizon can be constructed from these instantaneous prices. Next we use the instantaneous price corrections to study how concerns about robustness influence security prices.

### 7.3. Pricing without Robustness

Standard price theoretic arguments link prices to marginal rates of substitution. This link carries over to asset pricing and implies a direct connection between the marginal rate of substitution for the numeraire consumption good and the stochastic discount factor. Let  $\{mu(x_t)\}$  denote the logarithm of the marginal utility process for the numeraire consumption good. Then the familiar asset pricing formula is

$$\mathcal{P}_s \phi(x) = E(\phi(x_{t+s}) q_{t,t+s}^r | x_t = y)$$

where the stochastic discount factor is the intertemporal marginal rate of substitution:

$$q_{t,t+s}^r = \exp(-\delta s) \exp[mu(x_{t+s}) - mu(x_t)]. \quad (7.5)$$

For instance, in the power utility model

$$\frac{\exp[(1-\gamma)c]}{1-\gamma},$$

the (log) marginal utility is

$$mu(x) = -\gamma\psi(x)$$

where the robust solution for log consumption is  $\psi$ .

To relate this representation of pricing to (7.1), we use the stochastic integral representation of the evolution of the marginal utilities

$$mu(x_{t+s}) = mu(x_t) + \int_t^{t+s} \mu_m(x_u) du + \int_t^{t+s} \lambda_m(x_u) \cdot dB_u$$

where

$$\mu_m = \frac{\partial mu}{\partial x} \cdot \mu^* + \frac{1}{2} \text{trace} \left( \frac{\partial^2 mu}{\partial x \partial x'} \Sigma^* \right)$$

and

$$\lambda_m = \Lambda^{*'} \frac{\partial mu}{\partial x}.$$

Thus, the counterpart to relating prices to marginal rates of substitution is

$$\eta = \delta - \mu_m$$

$$\lambda = -\lambda_m,$$

with the risk-free rate given by  $\delta - \mu_m - \frac{1}{2} \lambda_m \cdot \lambda_m$ . These formulas match Breeden's formulas for the instantaneous risk free rate and the factor risk prices, though we have obtained them by a different argument than Breeden's. In the special case of the power utility model, the factor risk prices are

$$\begin{aligned} \lambda &= \lambda^* \\ &= \gamma \Lambda^{*'} \frac{\partial \psi}{\partial x} \end{aligned}$$

and hence depend linearly on the power parameter  $\lambda$ . Holding fixed the consumption allocation, larger values of  $\gamma$  (and hence more risk aversion) increase the magnitude of the risk prices.

#### 7.4. Incorporating Robustness

Our analysis of robust updating of value functions made us think of a preference for robustness as causing particular distortions of beliefs. These same distortions show us how to adjust intertemporal prices for robustness.

Let  $V^*$  denote the value function adjusted for robustness. For pricing we now form the stochastic discount factor as the product:

$$q_{t,t+s} = q_{t,t+s}^r q_{t,t+s}^u$$

where the term  $q_{t,t+s}^r$  is given by intertemporal marginal rate of substitution formula (7.5) and

$$q_{t,t+s}^u \equiv \frac{\exp\left[-\frac{V^*(x_{t+s})}{\theta}\right]}{E\left(\exp\left[-\frac{V^*(x_{t+s})}{\theta}\right] \mid x_t\right)}$$

is the adjustment for model uncertainty. Notice that this latter term is always strictly positive and that its conditional mean equals one for all  $s \geq 0$ . Moreover, the process  $\{q_{t,t+s}^u : s \geq 0\}$  is an exponential martingale and thus can be represented as:

$$q_{t,t+s}^u = \exp\left(\int_t^{t+s} \left[-\frac{1}{2}g^*(x_u) \cdot g^*(x_u) du + g^*(x_u) \cdot dB_u\right]\right)$$

where

$$g^*(x) = -\frac{1}{\theta}\Lambda^*(x)' \frac{\partial V^*(x)}{\partial x}.$$

Again the  $du$  component  $-\frac{1}{2}g^*(x_t) \cdot g^*(x_t)$  is the usual *log normal* adjustment needed for the martingale property to be satisfied. For this to be a valid representation, we require that the counterpart to Assumption 2 be satisfied.

To depict the robust pricing generator, we set

$$\begin{aligned} \rho &= \delta - \mu_m - \frac{1}{2}g^* \cdot g^* \\ \lambda &= -\lambda_m - g^*. \end{aligned}$$

In particular, the so called *factor risk prices* now have a model uncertainty component given by the drift distortion for the Brownian motion from the solution to the robust resource allocation problem. The implied risk-free rate is:

$$\rho^* = \left(\delta - \mu_m - \frac{1}{2}\lambda_m \cdot \lambda_m\right) - g^* \cdot \lambda_m \quad (7.6)$$

where the first term coincides with Breeden's formula and the second term  $-g^* \cdot \lambda_m$  adjusts interest rates to accommodate robustness.

<sup>23</sup> Hansen, Sargent, and Tallarini (1999) studied permanent income economies in which the riskfree rate is pinned down by the technology. They adjusted the subjective discount factor to compensate for the impact of formula (7.6) in equilibrium.

PROPOSITION 7.1. *Suppose that Assumption 7.1 and Assumption 7.2 are satisfied. The factor prices for the Brownian increments are given by  $-\lambda_m - g^*$  where  $-\lambda_m$  is the vector of risk prices (that ignore concerns about model misspecification) and  $g^*$  is the vector of model uncertainty prices. The vector of model uncertainty prices coincide with drift distortion for the Brownian motion from the robust resource allocation problem.*

The additive decomposition asserted in Proposition 7.1 was obtained as an approximation for linear-quadratic, Gaussian optimal resource allocation problems by Hansen, Sargent, and Tallarini (1999). By studying continuous time diffusion models we have been able to sharpen their results and relax the linear-quadratic specification of constraints and preferences.

### 7.5. An Example of Robust Pricing

Consider a special case of example two above. Let there be  $n$  technologies and assume

$$dk_t = \sum_{i=1}^n [(k_{it})^{\alpha_i} dy_{it} - \delta_i k_{it}] - c_t dt$$

with

$$dy_{it} = \mu_i dt + \sigma_i \cdot (g_t + dB_t).$$

Here the  $dy_{it}$  terms do not depend on exogenous variables. This economy can be modeled with a single state variable capital, denoted  $k_t$ , and  $n$  choice variables: one choice variable for consumption and  $n - 1$  choice variables for allocating capital to technologies. We apply the pricing formulas to this model.

The first order conditions include

$$U'(c) = \frac{\partial V}{\partial k}. \quad (7.7)$$

Note that

$$\frac{\partial mu}{\partial k} = \left( \frac{U''(c)}{U'(c)} \right) \frac{\partial c}{\partial k} \quad (7.8a)$$

$$g^* = -\frac{1}{\theta} \Lambda^* U'(c). \quad (7.8b)$$

Equation (7.8a) holds whenever the agents log marginal utility function is differentiable. Equation (7.8b) holds in this example because the state is one dimensional and first order condition (7.7) holds. It follows that the robust contribution to the risk-free rate is

$$\begin{aligned} -g^* \lambda_m &= -\frac{1}{\theta} \Sigma^* U''(c) \frac{\partial c}{\partial k} \\ &= -\frac{1}{\theta} \Sigma^* \frac{\partial^2 V}{\partial k^2}. \end{aligned}$$

The dependence of the effect of a preference for robustness on  $k$  is determined by the product of  $\Sigma^*$  and the second derivative of the value function with respect to capital. In general  $\Sigma^*$  depends on the optimal allocation of capital to technologies. In the special case of one technology,  $\Sigma^*$  depends only upon  $k$ .

Just as the magnitude of the factor risk prices are increasing in the curvature (risk aversion) parameter  $\gamma$ , the magnitude of the model uncertainty prices are increasing the robustness parameter  $\frac{1}{\theta}$ . Prior to studying in more detail how a concern about model misspecification alters security market prices, we use results from statistical decision theory to guide our choice of how much robustness is reasonable.

## 8. Detection Probabilities

This section uses statistical model discrimination to guide the choice of the robustness parameter  $\theta$ . We study a problem in which a decision-maker selects between two models, say model A and model B. For instance, model A might be the approximating model and model B the (constrained) worst case model from the two-player game. We then ask how much data are needed to tell these models apart. The aim is to eliminate models that are easy to tell apart statistically. It is not plausible to set the robustness parameter to be so small that we tailor decisions to be robust against alternatives that can be detected with high confidence with a limited amount of data. It would seem harder to justify making decisions robust to model departures that are readily discernible from data.

We study a statistical problem that is simpler than the one that decision makers confront in our economy. They face a large set of models. We consider statistical discrimination between only two Markov processes. It is easy to envision a much harder problem with a large set of models that makes probabilistic assignments difficult or impossible. Our use of a simplified problem to constrain the robustness parameter is in the same spirit as the use of highly stylized choice problems with prespecified probabilities to ascertain a reasonable amount of risk aversion to impose on decision-making.

In deriving our results, we develop an instantaneous measure of the statistical distance between models. This instantaneous measure will turn out to be closely related to the model uncertainty component of the equilibrium prices assigned to the Brownian motion increments. This close connection permits us to link the model uncertainty premium in security prices to the hypothetical statistical discrimination between models.

---

<sup>24</sup> See Pratt (1965) and Cochrane (1997).

### 8.1. A Simplified Decision Problem

We consider the case in which there are direct measurements on the state vector  $\{x_t : 0 \leq t \leq N\}$ . The aim is to discriminate between two Markov models: model A and model B. The two models are restricted to be absolutely continuous with respect to each other, which makes statistical discrimination difficult. Prior probabilities of one-half are imposed on each model, and a continuous record of data is observed. After an interval of time  $N$  elapses, a model is selected by maximizing the posterior probabilities.

Two types of errors are possible, choosing model A when B is correct and model B when model A is correct. We weight these errors by the prior probabilities and, following Chernoff (1952), we study the error probabilities as the sample interval becomes large. When the error probabilities decay slowly as a function of the sample interval, we say that the models are difficult to tell apart. To study this behavior, we are led to use a notion of entropy first introduced by Chernoff. As we will see, in the case of Markov diffusions, Chernoff entropy is closely related to the relative notion of entropy from our formulation of robust decision-making.

### 8.2. Chernoff entropy

Chernoff (1952) originally constructed his entropy measure for data that are independent and identically distributed. Let  $E$  denote the expectation operator associated with an approximating model A; let  $w$  denote a strictly positive function; and let  $w(y)/E(y)$  denote the density of a candidate model B *vis a vis* this approximating model. We divide by the expectation so that the resulting function is a valid density (has population mean one). An input into Chernoff's measure of entropy (indexed by  $\alpha$ ) is

$$H_\alpha = \frac{E(w^\alpha)}{[E(w)]^\alpha}$$

where  $0 < \alpha < 1$ . Since  $w^\alpha$  is a concave function, it follows that  $H_\alpha < 1$  and is equal to one when  $w$  is constant. Notice that when  $w$  is constant, the probability models A and B coincide. Chernoff entropy is defined to be

$$K_\alpha = -\log H_\alpha,$$

which is nonnegative and equal to zero when  $w$  is constant.

---

<sup>25</sup> A limiting version of Chernoff entropy is relative entropy:

$$\begin{aligned} I &= \lim_{\alpha \rightarrow 1} \frac{K_\alpha}{1 - \alpha} \\ &= E[w \log(w)] - \log(Ew) \end{aligned}$$

(*e.g.* see Vajida, 1990).

Chernoff (1952) used his entropy measure to deduce a bound on the error probabilities for Bayesian model selection. Adding a refinement due to Hellman and Raviv (1970), a bound for the probability  $\epsilon_N$  of making a classification error with sample size  $N$  is:

$$\epsilon_N \leq \inf_{0 \leq \alpha \leq 1} \frac{1}{2} \exp[-NK_\alpha(w)]. \quad (8.1)$$

Thus the probability of making a mistake decays at least exponentially as the sample size increases. To obtain the best possible decay rate bound, we are led to compute:

$$\kappa \equiv \sup_{0 \leq \alpha \leq 1} K_\alpha$$

When  $\kappa$  is large, it is easy to tell models apart, as indicated by the fast rate of exponential decay in the probabilities. In contrast, when  $\kappa$  is small, model discrimination can be problematic. This entropy measure of discrepancy is a distance measure: we would obtain the same positive number if we reversed roles of the two models.

Consider the following:

EXAMPLE 8.1. Suppose that under model A (say the approximating model) the data are normally distributed with mean  $\mu^*$  and nonsingular covariance matrix  $\Sigma^*$  and that under an alternative model B, the data are normally distributed with mean  $\hat{\mu}$  and covariance  $\Sigma^*$ . Then:

$$K_\alpha = \frac{(1-\alpha)\alpha}{2} (\hat{\mu} - \mu^*)' (\Sigma^*)^{-1} (\hat{\mu} - \mu^*).$$

The decay rate bound is:

$$\kappa = \frac{1}{8} (\hat{\mu} - \mu^*)' (\Sigma^*)^{-1} (\hat{\mu} - \mu^*). \quad (8.2)$$

While this example is very special, diffusion models lead to similar formulas because of the normality of the underlying Brownian motions.

### 8.3. Markov Formulation

Newman and Stuck (1979) deduced a Markov counterpart to the Chernoff bounds. A mathematical formulation like that used for pricing applies here. Take the model A (the approximating model) to be a Markov diffusion with drift coefficient  $\mu^*$  and diffusion matrix  $\Sigma^* = \Lambda^* \Lambda^{*'}$ . As in the construction of Markov perturbations, take the comparison model to be Markovian constructed by appending a state-dependent drift to the Brownian motion. Thus the altered drift for the diffusion can be depicted as

$$\hat{\mu} = \mu^* + \Lambda^* g \quad (8.3)$$

for some vector  $g$  of functions of the Markov state. As argued above, these departures preserve absolute continuity.

To study detection errors, we imitate our formulation of pricing by building a family (semigroup) of operators indexed by  $s$

$$\mathcal{D}_s^\alpha (\phi) (y) = E [\phi (x_{t+s}) d_{t,t+s}^\alpha | x_t = y]$$

where as in the pricing formulation the ‘stochastic discount factor’ satisfies (compare with equation (7.2)):

$$d_{t,t+s}^\alpha = \exp \left( - \int_t^{t+s} [\eta_\alpha (x_u) du + \lambda_\alpha (x_u) \cdot dB_u] \right).$$

This family of operators will be constructed so as to give the following (Chernoff) probability bound on the detection errors:

$$\epsilon_N = \frac{1}{2} \inf_{0 \leq \alpha \leq 1} \mathcal{D}_N^\alpha 1 \quad (8.4)$$

Newman and Struck (1979) provide us with a characterization of the semigroup  $\{D_s^\alpha : s \geq 0\}$  and the generator of that semigroup for each choice of  $\alpha$ . In particular,

$$\eta_\alpha = \frac{\alpha}{2} g' g$$

and

$$\lambda_\alpha = -\alpha g,$$

which is just  $\alpha$  times the drift distortion for the Brownian motion vector implicit in the construction of  $\hat{\mu}$  (see (8.3)).

Using formula (7.4) from the previous section, the generator  $\mathcal{C}^\alpha$  for the semigroup  $\{D_s^\alpha : s \geq 0\}$  is

$$\mathcal{C}^\alpha \phi = \mathcal{A}\phi + \alpha \Lambda^{*'} \frac{\partial \phi}{\partial x} - \left( \frac{\alpha}{2} g' g - \frac{\alpha^2}{2} g' g \right) \phi$$

As with pricing (see (7.3)), it is convenient to decompose the generator into two components:

$$\mathcal{C}^\alpha \phi = \left[ \mathcal{A}\phi + \alpha \frac{\partial \phi'}{\partial x} \Lambda^* g \right] - \left[ \frac{\alpha(1-\alpha)}{2} g'g \right] \phi$$

The first component distorts the expectations in the same manner as preferences for robustness, except with  $g^*$  replaced by  $\alpha g$ . The drift for the distortion in the conditional expectation operators is  $\mu + \alpha \Lambda^* g$ , while the diffusion covariance matrix remains the same.

The second component is of particular interest. The counterpart to the instantaneous risk free rate is now

$$\begin{aligned} -\lim_{s \downarrow 0} \frac{\log \mathcal{D}_s^\alpha}{s} &= -\mathcal{C}^\alpha 1 \\ &= \frac{\alpha(1-\alpha)}{2} g'g \\ &\equiv K'_\alpha(g) \end{aligned}$$

which is a simple quadratic form in the drift distortion of the Brownian motion. Localizing minimization problem (8.4), we are led to compute

$$\begin{aligned} \kappa' &= \max_{0 \leq \alpha \leq 1} K'_\alpha \\ &= \frac{1}{8} g'g \end{aligned}$$

which is typically state dependent. Thus the largest instantaneous discount rate for the detection-error probabilities is given by setting  $\alpha = \frac{1}{2}$  and the resulting rate is proportional to the entropy measure that enforces robustness. .

As we have seen, in the case of diffusions, differences between the Chernoff measure and conditional relative entropy agree up to a scale factor. The measure is symmetric in the sense that the roles of  $\mu^*$  and  $\hat{\mu}$  can be interchanged. The instantaneous rate,  $\kappa'$  is the diffusion analog to the constant rate  $\kappa$  in Example 1 (see (8.2)). The inverse covariance matrix is absent in the former expression because the vector of Brownian motions are standardized.

We will have more to say about the connection between pricing and detection in the next section of the paper.

<sup>26</sup> There will be a different solution for  $\alpha$  if we solve (8.4) for over a finite interval  $N$ .

<sup>27</sup> Typically, a long run measure of Chernoff entropy is deduced by locating the principal eigenvalue of the generator  $\mathcal{C}^\alpha$ .

<sup>28</sup> In classical statistics, conditional relative entropy can be used to deduce decay rates for type II errors holding type one errors fixed. These latter probabilities decay more quickly because of the decision to hold fixed the type I error independent of sample size. The faster decay shows up in the scale factor of  $1/2$  instead of  $1/8$  on  $g'g$  .

#### 8.4. Further Discussion

We indicated above that the statistical decision problem posed in this section is simple. It entails a pairwise comparisons of *ex ante* equally likely models. This feature is nice because it gives rise to a statistical measure of distance. The models under contention are both Markov models, which greatly simplifies the bounds on the probabilities of making a mistake when choosing between models. The implicit loss function needed to justify model choice based on the maximization of posterior probabilities is symmetric (*e.g.* see Chow, 1957). Finally, the decision-maker is compelled to select a specific model after a fixed amount of data are gathered.

Chernoff-type bounds can be obtained when there are more than two models. They are also known to be of value when the decision problem is extended to include the possibility of waiting for more data before making a decision. (*e.g.* see Hellman and Raviv, 1970). Like the problem we describe, these generalizations can be posed as Bayesian problems with explicit loss functions and prior probabilities.

However, model misspecification problems appear to us to be ones in which there is a rich (infinite dimensional) collection of potential models, including many that are considerably more complicated than our Markov benchmark. Nevertheless, the (constrained) worst case model remains Markov, and the tools just described can be employed to study a pairwise comparison between the approximating model and its robust counterpart. When there is a wide variety of competing models, the construction of prior probabilities is problematic and the computation of value functions over the entire set is impractical. On the other hand, robustness in decision-making aspires to asymmetries in the loss function. The aim is to construct decision rules that perform well over a range of models. Optimization for a given model may perform poorly if other models truly generate the data. This asymmetry works in the direction of making the robust rule more attractive than is the case in our simple statistical decision problem (with implicit symmetric losses to making classification errors). The simple problem we have studied is likely to make the decision-maker not choose the robust rule as often as he should.

While the statistical decision problem posed and analyzed here is by design too simple, it is useful in understanding or at least bounding how much robustness might be reasonable. We have, in effect, indexed robustness by the value of the parameter  $\theta$ . In assessing what a reasonable value is for  $\theta$ , we propose to check the magnitude of the worst case drift increment  $g^*$ . Associated with each value of the robustness parameter  $\theta$ , we can compute the Chernoff entropy process for the implied worst case. This entropy acts as an instantaneous discount rate for error bounds associated with a model misspecification. If

the departure is easy to detect (has small detection error probabilities), then the chosen  $\theta$  can be regarded as too small.

## 9. Market Price of Uncertainty

The market price of risk is typically defined to be the (absolute value of) the slope of the mean-standard deviation frontier for the full array of asset returns. Extensive empirical work has measured this tradeoff. One statement of the so called *equity premium puzzle* is that the observed risk return tradeoff implies that investors must be very risk averse (*e.g.* see Hansen and Jagannathan, 1991 and Cochrane and Hansen, 1992). In this section we show how a concern about misspecification changes the view of this puzzle. Using our decomposition of the Brownian increment prices, we attribute part of the steep slope to a concern about statistically small amounts of model misspecification.

### 9.1. Instantaneous Frontier

In characterizing the empirical implications, we analyze the implied tradeoff between risk and return using the drift and diffusion coefficients from the approximating model. Consider the gross return

$$r_{t+s,t} = \frac{\phi(x_{t+s})}{\mathcal{P}_s \phi(x_t)}$$

of an asset with payoff  $\phi(x_{t+s})$  at date  $t+s$ . The local mean or  $dt$  increment for this return is

$$\mu_r = \frac{\mathcal{A}\phi}{\phi} - \frac{\mathcal{B}\phi}{\phi}$$

where the first term on the right-hand side reflects local mean of the payoff (the numerator of the return) and the second term comes from the dependence of the pricing function on  $s$  (the denominator of the return). Substituting from formula (7.3), we can express the difference between the generator  $\mathcal{A}$  of the conditional expectations from the generator  $\mathcal{B}$  of the pricing operators as

$$\mu_r = \rho + \beta \cdot \lambda \tag{9.1}$$

where  $\rho$  is the instantaneous real interest rate,  $\lambda$  is the vector of factor prices for the Brownian increments and  $\beta$  is the vector of factor loadings for the Brownian increments of the instantaneous return. These loadings are given by

$$\beta = \Lambda^{*'} \frac{\partial \log \phi}{\partial x}.$$

Thus, once we know the vector  $\beta$  of factor loadings, the instantaneous mean return,  $\mu_r$ , is pinned down by the equilibrium pricing relation. The difference between the local mean of the return and the riskfree rate is the dot product between the factor loading and the factor price vector.

Since the standard deviation of the asset is  $|\beta|$ , by the Cauchy-Schwarz inequality we may bound the Sharpe ratio (the excess return over the standard deviation) via

$$\frac{|\mu_r - \rho|}{|\beta|} = \frac{\beta \cdot \lambda}{|\beta|} \leq |\lambda|$$

Notice that the inequality is attained by setting  $\beta = \lambda$ . Thus  $|\lambda|$  is the slope of the risk-return tradeoff computed with the means and standard deviations from the approximating model.

## 9.2. Decomposing the slope

Usually, the slope of the frontier is used partly to gauge how risk averse investors are. A big slope suggests that investors participating in security markets are highly risk averse. For instance, this can be seen from the power utility model without robustness since the slope is

$$|\lambda_m| = \gamma \left| \Lambda^{*'} \frac{\partial \psi}{\partial x} \right|$$

where  $\psi$  is the log consumption function. Thus the magnitude of the  $\lambda$  is dictated in part by risk aversion. Recall, however, that Brownian increment prices for the economy with robustness can be decomposed as

$$\lambda = -\lambda_m - g^*$$

where  $g^*$  is governed by how robust agents are in their decision making. Thus we are led to define the market price of risk to be

$$\text{mpr} = |\lambda_m|,$$

while the market price of model uncertainty is

$$\begin{aligned} \text{mpu} &= |g^*| \\ &= \frac{1}{\theta} \left| \Lambda^{*'} \frac{\partial V}{\partial x} \right|. \end{aligned}$$

The slope of the frontier of the mean-standard deviation frontier reflects both components. We next consider what are reasonable magnitudes for the market price of uncertainty.

mpu	Chernoff	probability bound	probability
.05	.00031	.470	.362
.10	.00125	.389	.240
.20	.00500	.184	.078
.40	.02000	.009	.000

**Table 9.1:** Prices of Model Uncertainty and Detection-Error Probabilities,  $N = 200$ .

### 9.3. Entropy and market price of uncertainty

As we have just defined it, the market price of uncertainty is closely tied to the Chernoff entropy measure  $\kappa'$  for the the endogenously chosen worst case ( $g = g^*$ ) by the formula

$$\text{mpu} = \sqrt{8\kappa'}$$

Thus the market price of uncertainty is directly tied to the instantaneous detection error rate for discriminating between the approximating model and the (constrained) worst case model. We use this formula to study detection probabilities associated with alternative magnitudes of  $\theta$  and hence alternative magnitudes of market price of uncertainty.

Table 1 reports values of mpu, Chernoff entropy, the probability-error bound (8.1), and the actual probability of detection  $\epsilon_N$  on the left side of (8.1) (which we can calculate analytically in this case) for two values of  $N$ . The probability bounds and the probabilities are computed under the simplifying assumption that the Chernoff entropy are state independent. With constant drift and diffusion coefficients, the log-likelihood ratio is normally distributed, which allowed us easily to compute the actual detection-error probabilities.

We took the sample interval to 200 thinking of a quarter as corresponding to a unit of time. Alternatively, we might have used and a sample interval of 600 to link to monthly postwar data. The market prices of risk and uncertainty are associated with specific time unit normalizations. Since, at least locally, drifts coefficients and diffusion matrices scale linearly with the time unit, the market prices of risk and uncertainty scale with the square root of the time unit.

The numbers in Table 1 indicate how market prices of uncertainty somewhat less than .2 are associated with associated with misspecified models that are hard to detect. However, market prices of uncertainty of .40 are associated with easily detectable alternative models.

The tables also reveal that although the probability bounds are weak, they display patterns similar to those of the actual probabilities.

Empirical measures of the slope of the mean-standard deviation frontier are about .25 for quarterly data. Given the absence of volatility in aggregate consumption, risk considerations only explain a small component of the measured risk-return tradeoff using aggregate data (Hansen-Jagannathan, 1991). In contrast, our calculations suggest that concerns about statistically small amounts of model misspecification could account for a substantial component of the empirical measures.

So far our calculations have assumed that Chernoff entropy is independent of the Markov state. We also computed detection-error bounds and detection-error probabilities for the robust permanent income model of HST (Hansen, Sargent, and Tallarini, 1999). The Chernoff entropies are state dependent for this model, but the probability bounds can still be computed numerically. We input the parameter values from HST's discrete-time model to form an approximate continuous-time robust permanent income model using conversions in MATLAB control toolbox. In their economy and in its continuous-time counterpart, consumption and investment profiles along with real interest rates remain fixed when we change the robustness parameter. This is accomplished by simultaneously altering the subjective discount rate to offset the precautionary motive for saving. It can be shown that the worst case  $g^*$  vector is proportional to the marginal utility of consumption and hence is highly persistent. This is to be expected because the permanent income model is well designed for transient fluctuations, but is vulnerable to model misspecifications that are highly persistent under the approximating model. Under the approximating model the marginal utility process is a martingale, but under the (constrained) worst case this process is an explosive scalar autoregression. The choice of  $\theta$  dictates the magnitude of the autoregressive coefficient for the marginal utility process. We now study the corresponding detection-error probabilities. In this case, the Chernoff entropy is state dependent, leading us to compute the solution to problem (8.1) numerically.

Notice that the results in Table 2 are close to those of Table 1, although the detection-error probabilities are a little bit lower in Table 1. Recall that the Table 1 results abstract from state dependence. As shown by HST, the fluctuations in  $|g^*|$  are quite small relative to its level, which may help to explain the similar results. It still is the case that statistical detection is difficult for market prices of uncertainty up to about say half that of the empirical measure. Other model complications are also required to account fully for the steep slope of the frontier. For instance, the model could be extended to include the

$\frac{1}{\theta}$	mpu	probability bound	simulated probability
.00005	.058	.456	.338
.00010	.116	.332	.181
.00015	.174	.178	.060

**Table 9.2:** Prices of Model Uncertainty and Detection-Error Probabilities for the Permanent Income Model,  $N = 200$ .

external habit specification of Campbell and Cochrane, 1999 and market frictions along the lines of Constantinides and Duffie, 1996, or Heaton and Lucas, 1996.

## 10. Conclusion

We have studied a duality relation between risk and robustness that emerges when we make a risk-sensitivity adjustment to continuation utilities, in the fashion of Epstein and Zin (1989). We use the duality to reinterpret the risk adjustment as expressing a preference for robustness. It quantifies a relationship between model misspecification and utility loss.

Our altered decision problem introduces a robustness parameter. In assessing reasonable amount of risk aversion in stochastic, general equilibrium models, it is common to explore preferences for hypothetical gambles in the manner of Pratt (1965). In assessing reasonable preferences for robustness, we propose using large sample detection probabilities for a hypothetical model selection problem. We envision a decision-maker as choosing to be robust to departures that are hard to detect statistically.

Associated with the duality result in decision-making is a duality result in the decentralized prices of securities. For simplicity, we reconsider a familiar asset pricing environment where shadow prices are inferred for increments of a multivariate Brownian motion. We

---

<sup>29</sup> We find it fruitful to explore concern about model uncertainty because these other model modifications are themselves only partially successful. To account fully for the market price of risk, Campbell and Cochrane (1999) adopt specifications with an arguably large amount of risk aversion during recessions. Constantinides and Duffie (1996) accommodate fully the high market prices of risk by attributing empirically implausible consumption volatility to individual consumers (see Cochrane, 1998). Finally, Heaton and Lucas (1998) show that reasonable amounts of proportional transaction costs can only explain about half of the equity premium puzzle.

show that factor risk prices have components that can be viewed as factor prices of model uncertainty. Thus what appears in security market data to be a large ‘risk premium’ may in part be a ‘model misspecification’ premium.

We see three important extensions to our current investigation. Like rational expectations models, we have side-stepped the issue of how decision-makers select an approximating model. Following the literature on robust control, we envision this approximating model to be analytically tractable, yet not to provide a correct model of the evolution of the state vector. The misspecifications we have in mind are small in a statistical sense but can otherwise be quite diverse. On the other hand, we have not formally modeled how agents learned the approximating model and why they do not bother to learn about potentially complicated misspecifications of that model. Incorporating forms of learning would be an important extension of our work.

The equilibrium calculations in our model currently exploit the representative agent paradigm in an essential way. Reinterpreting our calculations as applying to a multiple agent setting is straightforward in some circumstances (see Tallarini, 1998), but in general, even under rich market security market structures, multiple agent versions of this environment look fundamentally different from their single-agent counterparts (see Anderson, 1998). Thus, the introduction of heterogeneous decision-makers might result in new insights about the role of robustness in decision making.

Finally, we obtain our most complete characterizations of decision problems and the decentralized prices in the case of diffusion models. While the diffusion setup leads to some pedagogically useful simplifications, robustness considerations are likely to be particularly important in environments with large shocks that occur infrequently, as in Poisson models.

## 11. References

- Anderson, E.W. (1998), “Uncertainty and the Dynamics of Pareto Optimal Allocations,” University of Chicago Dissertation.
- Breeden, D. T. (1979), “An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities,” *Journal of Financial Economics* 7, 265-296.
- Brock, W. A. (1979), “An Integration of Stochastic Growth Theory and the Theory of Finance, Part I: The Growth Model,” *General Equilibrium, Growth and Trade*.
- Brock, W. A. and M. J. P. Magill (1979), “Dynamics Under Uncertainty,” *Econometrica* 47, 843-868.
- Brock, W.A. and L.J. Mirman (1972), “Optimal Growth Under Uncertainty: The Case of Discounting,” *Journal of Economic Theory*, 4:3, 479-513.

- Campbell, J. H. and J. H. Cochrane (1999), "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107:2, 205-251.
- Chernoff, H. (1952), "A Measure of Asymptotic Efficiency for Tests of a Hypothesis Based on the Sum of Observations," *Annals of Mathematical Statistics* 23, 493-507.
- Cochrane, J. H. (1997), "Where is the Market Going? Uncertain Facts and Novel Theories," *Economic Perspectives*, 21:6, 3-37.
- Cochrane, J. H. and L. P. Hansen (1992), "Asset Pricing Lessons for Macroeconomics," 1992 NBER Macroeconomics Annual, O. Blanchard and S. Fischer, (eds), 1115-1165.
- Constantinides, G. M. and D. Duffie (1996), "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 104, 219-240.
- Cox, J. C., J. E. Ingersoll, Jr. and S. A. Ross (1985), "An Intertemporal General Equilibrium Model of Asset Prices," *Econometrica*, 53, 363-384.
- Csiszar, I. (1991), "Why Least Squares and Maximum Entropy? An Axiomatic approach to Inference for Linear Inverse Problems," *The Annals of Statistics* 19:4, 2032-2066.
- Dow, J. and S. R. C. Werlang (1994), "Learning under Knightian Uncertainty: The Law of Large Numbers for Non-additive Probabilities, manuscript
- Durrett, R., (1996), *Stochastic Calculus: A Practical Introduction*, Boca Raton: CRC Press.
- Duffie, D. and L.G. Epstein (1992), "Stochastic Differential Utility," *Econometrica* 60:2, 353-394.
- Duffie, D. and P.L. Lions (1992), "PDE Solutions of Stochastic Differential Utility," *Journal of Mathematical Economics* 21, 577-606.
- Dupuis, P. and R. S. Ellis (1997), *A Weak Convergence Approach to the Theory of Large Deviations*, New York: John Wiley and Sons.
- Elliot, R. J., L. Aggoun, and J. B. Moore (1995), *Hidden Markov Models: Estimation and Control*, Springer: New York.
- Ellsberg, D. (1961), *Quarterly Journal of Economics*.
- Epstein, L.G. and T. Wang (1994), "Intertemporal Asset Pricing Under Knightian Uncertainty," *Econometrica* 62:3, 283-322.
- Epstein, L.G. and S.E. Zin (1989), "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica* 57:4, 937-969.
- Ethier, S.N. and T.G. Kurtz (1986), *Markov Processes: Characterization and Convergence*, NY: John Wiley and Sons.
- Fleming, W.H. and W.M. McEneaney (1995), "Risk-Sensitive Control on an Infinite Time Horizon," *SIAM Journal of Control and Optimization* 33:6, 1881-1915.
- Gilboa, I. and D. Schmeidler (1989), "Maxmin Expected Utility with Non-unique Prior," *Journal of Mathematical Economics* 18, 141-153.

- Hansen, L.P. and R. Jagannathan (1991), "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 99:2, 225-262.
- Hansen, L.P. and T.J. Sargent (1993), "Seasonality and Approximation Errors in Rational Expectations Models," *Journal of Political Economy*, 55, 21-55
- Hansen, L.P. and T.J. Sargent (1995), "Discounted Linear Exponential Quadratic Gaussian Control," *IEEE Transactions on Automatic Control* 40:5, 968-971.
- Hansen, L.P. and T.J. Sargent (1999), "Discounted Robust Linear Decision and Estimation in the Frequency Domain," manuscript.
- Hansen, L.P., T.J. Sargent and T.D. Tallarini, Jr. (1999), "Robust Permanent Income and Pricing," forthcoming in *Review of Economic Studies*.
- Heaton, J. C. and D. J. Lucas (1996), "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing," *Journal of Political Economy*, 104, 443-487.
- Hellman, M. E. and J. Raviv, (1970), "Probability Error, Equivocation, and the Chernoff Bound," *IEEE transactions on Information Theory*, volume 16, 368-372.
- James, M.R. (1992), "Asymptotic Analysis of Nonlinear Stochastic Risk-Sensitive Control and Differential Games," *Mathematics of Control, Signals, and Systems* 5, 401-417.
- Karatzas, I. and S. E. Shreve (1991), *Brownian Motion and Stochastic Calculus*, Second Edition, New York: Springer-Verlag.
- Kurz, M. (1997), *Endogenous Economic Fluctuations: Studies in the Theory of Rational Beliefs*, New York: Springer.
- Lucas, R. E., Jr. (1978), "Asset Prices in an Exchange Economy," *Econometrica*, 46:6, 1429-1445.
- Lucas, R.E., Jr. and E.C. Prescott (1971), "Investment under Uncertainty," *Econometrica*, 39:5, 659-681.
- Lucas, R.E. , Jr (1976), "Econometric Policy Evaluation: A Critique," *Journal of Monetary Economics*, 1:2, 19-46.
- Merton, R. (1973), "An Intertemporal Capital Asset Pricing Model," *Econometrica* 41, 867-888.
- Newman, C. M. and B. W. Stuck (1979), "Chernoff Bounds for Discriminating Between Two Markov Processes," *Stochastics*, 2, 139-153.
- Pratt, John W. (1965), "Risk Aversion in the Small and in the Large," *Econometrica*, 32(1-2), 122-136
- Runolfsson, T. (1994), "The Equivalence between Infinite-horizon Optimal Control of Stochastic Systems with Exponential-of-integral Performance Index and Stochastic Differential Games," *IEEE Transactions on Automatic Control* 39, 1551-1563.
- Runolfsson, T. (1995), "Risk-sensitive Control of Stochastic Hybrid Systems," *Proceedings of the 34th Conference on Decision and Control*, New Orleans, LA.

Shore, J.E. and R.W. Johnson (1980), "Axiomatic Derivation of the Principle of Maximum Entropy and the Principle of Minimum-Cross Entropy," *IEEE Transactions on Information Theory* 26:1, 26-37.

Sims, C. A. (1993), "Rational Expectations Modeling with Seasonally Adjusted Data," *Journal of Econometrics* 55.

Tallarini, Jr., T.D. (1998), "Risk-Sensitive Real Business Cycles," forthcoming in *Journal of Monetary Economics*.

Vajda, I. (1990), "Generalization of Discrimination-Rate Theorems of Chernoff and Stein," *Kybernetika*, 26, 273-287.

Weil, P. (1993), "Precautionary Savings and the Permanent Income Hypothesis," *Review of Economic Studies* 60, 367-383.