

Robustness and Uncertainty Aversion

Lars Peter Hansen Thomas J. Sargent
Gauhar A. Turmuhambetova Noah Williams*

February 25, 2002

Abstract

Max-min expected utility theory uses multiple prior distributions to represent uncertainty aversion. Robust control models a decision maker who fears that the data are generated by an unknown perturbation of his approximating model. We link the two approaches by interpreting the set of perturbations in robust control theory as the multiple priors of the max-min expected utility theory. We use a Brownian motion information structure and construct recursive versions of max-min expected utility theory and robust control theory.

1 Introduction

This paper links the max-min expected utility theory of Gilboa and Schmeidler (1989) to applications of stochastic robust control theory by James (1992), Petersen, James, and Dupuis (2000), and Anderson, Hansen, and Sargent (2000). The max-min expected utility theory represents uncertainty aversion with preference orderings over stochastic processes of decisions c and states x , for example, of the form

$$\inf_{q \in Q^*} E_q \left[\int_0^\infty \exp(-\delta t) U(c_t, x_t) dt \right] \quad (1)$$

where Q^* is a convex set of probability measures over (c, x) and δ is a discount rate. We shall call c consumption and include the state vector x in U to accommodate time nonseparabilities. In Gilboa and Schmeidler's theory, minimization over Q^* represents aversion to uncertainty. Gilboa and Schmeidler's theory leaves open how to specify Q^* in particular applications.

Criteria like (1) also appear as objective functions in robust control theory, where minimization over Q^* is a device for promoting robustness to misspecification of the decision

*We thank Fernando Alvarez, Gary Chamberlain, Ivar Ekeland, Peter Klibanoff, Michael Allen Rierson, Aldo Rustichini, Jose Scheinkman, Christopher Sims, and especially Costis Skiadis for valuable comments on an earlier draft. Sherwin Rosen encouraged us to write this paper.

maker's single explicitly specified approximating model, which lies in the interior of Q^* . Robust control theory generates Q^* by statistically perturbing that approximating model. By contemplating perturbations, the decision maker acknowledges possible model misspecification. Robust control theory represents Q^* implicitly through a positive parameter θ that penalizes relative entropy, a measure of model misspecification, thereby capturing the idea that the decision maker's model is a *good* approximation. This penalty parameter is used to formulate a 'multiplier problem' related to but distinct from (1). This paper describes how to transform that multiplier problem into a 'constraint problem' like (1). The constraint and multiplier problems differ in subtle ways, but the Lagrange multiplier theorem (Luenberger, 1969, pp. 216-221) connects them. The two problems imply different preference orderings over $\{c_t\}$ that nevertheless lead to the same decisions. We describe the senses in which the multiplier and constraint problems are both recursive, and therefore how both are time consistent. To facilitate comparisons to Anderson, Hansen, and Sargent (2000) and Chen and Epstein (2000), we cast our discussion within continuous-time diffusion models. The Brownian motion information structure used in this paper simplifies the solution to the stochastic version of the robust control problem.

Solutions of both problems satisfy Bellman equations. However, to make the constraint problem recursive requires augmenting the *state* to include a continuation value for relative entropy and also augmenting the control set of the minimizing agent to include an increment to the continuation value of entropy. Because it involves fewer states and controls, the multiplier problem is easier to solve.

We also discuss and defend the constraints on the allocation of relative entropy over time that are implicit in the recursive version of the constraint problem. Our recursive formulation makes the minimizing agent 'let bygones be bygones' by requiring that at time t he explore only misspecifications that allocate continuation entropy across yet to be realized events. We argue that this is a reasonable way of restricting the class of misspecifications that should concern the decision maker as time unfolds.

It can be useful to study decision rules from alternative vantage points. With this in mind, we describe two more interpretations of a robust control law. First, under a Bellman-Isaacs condition, there is a particular probability specification for the Brownian motion under which the robust control law is optimal in a Bayesian sense. A *pessimistic* twisting of the probability distribution associated with the approximating model supports this Bayesian interpretation. If the approximating model were actually true, this Bayesian decision maker would not have rational expectations.

A second interpretation maintains the Brownian motion model and rational expectations under the approximating model, but makes the preferences of the decision-maker be more *risk sensitive*. Risk-sensitive control theory, as initiated by Jacobson (1973), provides a tractable way to let decision rules be more responsive to risk through the use of an exponential adjustment to the instantaneous return function of the decision-maker. A variety of results in the control theory literature link risk-sensitivity to a concern about robustness, *e.g.* see James (1992). Hansen and Sargent (1995) and Anderson, Hansen, and Sargent (2000) formulate a risk-sensitive objective by using the recursive utility theory of Epstein and Zin (1989), Duffie

and Epstein (1992) and Schroder and Skiadis (1999) with an exponential risk adjustment to continuation values. Using recursive utility theory, Anderson, Hansen, and Sargent (2000) re-establish the robustness/risk-sensitivity connection for continuous-time Markov process models.

The remainder of this paper is organized as follows. Section 2 describes a standard stochastic control problem without concern about model misspecification. Section 3 defines and compares ‘time-zero’ versions of the multiplier and constraint robust control problems, each of which is posed as a zero-sum two-person game under mutual commitment to stochastic decision processes at time zero by an initial decision maker and a malevolent nature. To measure model misspecifications, section 3 defines a relative entropy for stochastic processes, and section 4 obtains representations for stochastic perturbations that are useful for solving commitment and recursive robust control problems. Section 5 then shows how the multiplier problem can be solved recursively. There we state a Bellman-Isaacs condition under which there is tight connection between the Markov perfect equilibrium and the date zero commitment equilibrium. Section 6 shows how to deduce the probability specification which render robust control processes optimal in a Bayesian sense. In section 7, we discuss the connections between our robust control problems and a particular form of recursive utility, known as risk sensitivity. Section 8 returns to the constraint problem, and shows how it can be solved recursively by augmenting the state and control vectors appropriately. Section 9 briefly compares and discusses the Bellman equations for the multiplier problem, the constraint problem, and the risk sensitive control problem. Section 10 defines and compares two preference relations over consumption sequences that are inspired by the multiplier and constraint robust control problems. Section 11 describes how both of these preference relations can be represented recursively. Section 12 concludes.

2 A Benchmark Resource Allocation Problem

As a benchmark for the rest of the paper, this section poses a discounted, infinite time optimal resource allocation problem in which the decision maker knows the model, and so has no concern about robustness to model misspecification. Later sections of the paper will assume that the decision maker regards his model as an approximation to some nearby unknown model that actually governs the data.

Let $\{B_t : t \geq 0\}$ denote a d -dimensional, standard Brownian motion on an underlying probability space (Ω, \mathcal{F}, P) . Let $\{\mathcal{F}_t : t \geq 0\}$ denote the completion of the filtration generated by this Brownian motion. For any stochastic process $\{g_t : t \geq 0\}$, we use g or $\{g_t\}$ to denote the process and g_t to denote the time t -component of that process. The actions of the decision-maker form a progressively measurable stochastic process $\{c_t : t \geq 0\}$, which means that the time t component c_t is \mathcal{F}_t measurable.¹ Let U be an instantaneous utility function

¹Progressive measurability requires that we view $c \doteq \{c_t : t \geq 0\}$ as a function of (t, ω) . For any $t \geq 0$, $c : [0, t] \times \Omega$ must be measurable with respect to $\mathcal{B}_t \times \mathcal{F}_t$ where \mathcal{B}_t is a collection of Borel subsets of $[0, t]$. See Karatzas and Shreve (1991) pages 4 and 5 for a discussion.

and C be the set of admissible control processes.

Definition 2.1. *Our benchmark control problem is:*

$$\sup_{c \in C} E \left[\int_0^\infty \exp(-\delta t) U(c_t, x_t) dt \right] \quad (2)$$

subject to:

$$dx_t = \mu(c_t, x_t)dt + \sigma(c_t, x_t)dB_t \quad (3)$$

where x_0 is a given initial condition.

Later we alter the Brownian specification for B to allow for model uncertainty.

We restrict μ and σ so that any progressively measurable control c in C implies a progressively measurable state vector process x . We maintain

Assumption 2.2.

$$\sup_{c \in C} E \left[\int_0^\infty \exp(-\delta t) U(c_t, x_t) dt \right]$$

subject to (3) is finite.

Thus the objective for the control problem without model uncertainty has a finite upper bound.

We express a concern for robustness in terms of a family of stochastic perturbations to the Brownian motion process and measure these perturbations with *relative entropy*.

3 Robust Control Problems

In this section we pose two robust control problems and describe their relation. One problem is more convenient for computation while the other expresses an intertemporal version of preferences that conform to axioms of Gilboa and Schmeidler (1989) for portraying uncertainty aversion. We begin with a setting in which the maximizing and minimizing decision makers both choose once and for all at date zero.

Just as a random variable induces a probability distribution over the real line, the d -dimensional Brownian motion induces a distribution on a *canonical space* $(\Omega^*, \mathcal{F}^*)$ defined as follows. Let Ω^* be the space of continuous functions $f : [0, +\infty) \rightarrow \mathbb{R}^d$. Let \mathcal{F}_t^* be the Borel sigma algebra for the restriction of the continuous functions f to $[0, t]$, where open sets are defined using the sup-norm over this interval. Notice in particular that $\iota_s(f) = f(s)$ is \mathcal{F}_t^* measurable for each $0 \leq s \leq t$. Let \mathcal{F}^* be the smallest sigma algebra containing \mathcal{F}_t^* for $t \geq 0$. An event in \mathcal{F}_t^* restricts the properties of the continuous functions only on the finite interval $[0, t]$. The Brownian motion B induces a multivariate Wiener measure on $(\Omega^*, \mathcal{F}^*)$, which we denote q^0 .

3.1 Absolute Continuity

We are interested in probability distributions that are absolutely continuous with respect to Wiener measure. For any probability measure q on $(\Omega^*, \mathcal{F}^*)$, we let q_t denote the restriction to \mathcal{F}_t^* . In particular, q_t^0 is the multivariate Wiener measure over the event \mathcal{F}_t^* .

Definition 3.1. *A distribution q is said to be **absolutely continuous over finite intervals** with respect to q^0 if q_t is absolutely continuous with respect to q_t^0 for all t .²*

Let Q be the set of all distributions that are absolutely continuous with respect to q^0 over finite intervals.

Remark 3.2. *To capture the idea that they are difficult to detect from samples of finite length, we require that perturbations to an approximating model be absolutely continuous with respect to it over finite intervals. Finite interval absolute continuity is weaker than absolute continuity. When q is absolutely continuous with respect to q^0 over finite intervals, a strong Law of Large Numbers that applies to a process constructed with shocks $d\hat{B}_t$ governed by q^0 would not necessarily also apply if the shock process were governed by q instead. Time series averages that converge almost surely under q^0 may not converge under q , so that q can be distinguished from q^0 given an infinite amount of data.³ Absolute continuity over finite intervals, however, is sufficient to allow us to construct likelihood ratios between models for finite histories at any calendar date t .*

3.2 Relative Entropy

To limit the alternative models that the decision-maker entertains, we now construct a relative entropy measure for a perturbed stochastic process. Form $\tilde{\Omega} = \Omega^* \times \mathbb{R}^+$ where \mathbb{R}^+ is the nonnegative real line. Form the corresponding sigma algebra $\tilde{\mathcal{F}}$ as the smallest sigma algebra containing $\mathcal{F}_t^* \otimes \mathcal{B}_t$ for any t where \mathcal{B}_t is the collection of Borel sets in $[0, t]$; and form \tilde{q} as the product measure q with an exponential distribution with density $\delta \exp(-\delta t)$ for any $q \in Q$.

Consider a progressively measurable family $\phi \doteq \{\phi_t : t \geq 0\}$ on $(\Omega^*, \mathcal{F}^*)$. The \tilde{q} expectation of ϕ is by construction

$$\int \phi d\tilde{q} = \delta \int_0^\infty \exp(-\delta t) \int \phi_t dq_t dt$$

which is an exponential average.

²(Kabanov, Lipcer, and Sirjaev 1979) refer to this concept as *local* absolute continuity. Although Kabanov, Lipcer, and Sirjaev (1979) define local absolute continuity through the use of stopping times, they argue that their definition is equivalent to this “simpler one”.

³Our specification allows Q measures to put different probabilities on tail events, which prevents the measures from merging as Blackwell and Dubins (1962) show will occur under absolute continuity. See Kalai and Lerner (1993) and Jackson, Kalai, and Smordoninsky (1999) for implications of absolute continuity for learning.

We measure the discrepancy between the probability distributions q^0 and q as the *relative entropy* between \tilde{q} and \tilde{q}^0 :⁴

$$\mathcal{R}^*(q) = \delta \int_0^\infty \exp(-\delta t) \left(\int \log \left(\frac{dq_t}{dq_t^0} \right) dq_t \right) dt \quad (4)$$

Here $\frac{dq_t}{dq_t^0}$ is the Radon-Nikodym derivative of q_t with respect to q_t^0 .

Lemma 3.3. \mathcal{R}^* is convex on Q .

Proof. Relative entropy is known to be convex in the product measure \tilde{q} . Consider two measures \tilde{q}^1 and \tilde{q}^2 formed from the product of q^1 and q^2 with the exponential distribution with parameter δ . Then a convex combination of \tilde{q}^1 and \tilde{q}^2 is given by the product of the corresponding convex combination of q^1 and q^2 with the same exponential distribution. \square

3.3 Two Robust Control Problems

We now have the vocabulary to state two related robust control problems. We use Q as a family of distortions to the probability distribution. Initially we state the robust control problems in terms of the induced distributions. This facilitates Lagrange multiplier methods. Given a progressively measurable control c we solve the stochastic differential equation (3) to obtain a progressively measurable utility process

$$u(c_t, x_t) = v_t(c, B)$$

where $v(c, \cdot)$ is a progressively measurable family in $(\Omega^*, \mathcal{F}^*)$. In what follows we will drop the second argument of v_t when we integrate using induced distributions. With this notation, the objective can be represented as:

$$\int_0^\infty \exp(-\delta t) \left(\int v_t(c) dq_t \right) dt.$$

Notice that this representation incorporates the state evolution. Our benchmark control problem uses q^0 for q .

We consider two robust control problems:

Definition 3.4. A **multiplier** robust control problem is:

$$\tilde{J}(\theta) = \sup_{c \in C} \inf_{q \in Q} \int_0^\infty \exp(-\delta t) \left(\int v_t(c) dq_t \right) dt + \theta \mathcal{R}^*(q). \quad (5)$$

Definition 3.5. A **constraint** robust control problem is:

$$J^*(\eta) = \sup_{c \in C} \inf_{q \in Q(\eta)} \int_0^\infty \exp(-\delta t) \left(\int v_t(c) dq_t \right) dt \quad (6)$$

where $Q(\eta) = \{q \in Q : \mathcal{R}^*(q) \leq \eta\}$.

⁴Other measures of relative entropy for stochastic processes occur in the literature on large deviations. For example, see Dupuis and Ellis (1997) page 299.

The first problem is a stochastic counterpart to ones found in the robust control theory literature. The second embodies a date-zero version of the multiple priors model advocated by Gilboa and Schmeidler (1989), and is analogous to (1) if we let $Q^* = Q(\eta)$.

3.4 Relation Between Problems

As is typical in penalty formulations of decision problems, we can interpret the robustness parameter θ in the first problem as a Lagrange multiplier on the specification-error constraint $\mathcal{R}^*(q) \leq \eta$. This connection is regarded as self-evident throughout the literature on robust control and has been explored in the context of a linear-quadratic control problem, informally by Hansen, Sargent, and Tallarini (1999), and formally by Hansen and Sargent (2001a). Here we study this connection within our continuous time stochastic setting, relying heavily on developments in Petersen, James, and Dupuis (2000) and Luenberger (1969).

As a consequence of Assumption 2.2, the optimized objectives for the multiplier and constraint robust control problems must both be less than $+\infty$. These objectives could be $-\infty$, depending on the magnitudes of θ and η .

We use θ to index a family of multiplier robust control problems and η to index a family of constraint robust control problems. We call *admissible* only those nonnegative values of θ for which it is feasible to make the objective function greater than $-\infty$. If $\hat{\theta}$ is admissible, values of θ larger than $\hat{\theta}$ are also admissible, since these values only make the objective larger. Let $\underline{\theta}$ denote the greatest lower bound for admissible values of θ .

Given an $\eta > 0$, add $-\theta\eta$ to the objective in (5). For a given value of θ this has no impact on the control law.⁵ We motivate this subtraction by the Lagrange multiplier theorem (see Luenberger (1969, pp. 216-221)) and use the maximized value of θ to relate the multiplier robust control problem to the constraint robust control problem.

For a given c , the objective of the constraint robust control problem is linear in q and the entropy measure \mathcal{R}^* in the constraint is convex in q . Moreover, the family of admissible probability distributions Q is itself convex. We formulate the constraint version of the robust control problem as a Lagrangian:

$$\sup_{c \in C} \inf_{q \in Q} \sup_{\theta \geq 0} \int_0^\infty \exp(-\delta t) \left(\int v_t(c) dq_t \right) dt + \theta [\mathcal{R}^*(q) - \eta]$$

It is well known that the optimizing multiplier θ is degenerate for many choices of q . It is infinite if q violates the constraint and zero if the constraint is slack. We can exchange the order of the \sup_θ and \inf_q and still support the same value of q . The Lagrange Multiplier Theorem allows us to study:

$$\sup_{c \in C} \sup_{\theta \geq \underline{\theta}} \inf_{q \in Q} \int_0^\infty \exp(-\delta t) \left(\int v_t(c) dq_t \right) dt + \theta [\mathcal{R}^*(q) - \eta] \quad (7)$$

Unfortunately, the maximizing θ in (7) depends on the choice of c . In solving a robust control problem, we are most interested in the c that solves the constraint robust control

⁵However, it will alter which θ results in the highest objective.

problem. We can find the appropriate choice of θ by changing the order of \sup_c and \sup_θ to obtain:

$$\sup_{\theta \geq \underline{\theta}} \sup_{c \in C} \inf_{q \in Q} \int_0^\infty \exp(-\delta t) \left(\int v_t(c) dq_t \right) dt + \theta [\mathcal{R}^*(q) - \eta]$$

Suppose that \sup_θ is attained and call the optimizing value θ^* . When we fix θ at θ^* we are led to solve

$$\sup_{c \in C} \inf_{q \in Q} \int_0^\infty \exp(-\delta t) \left(\int v_t(c) dq_t \right) dt + \theta^* \mathcal{R}^*(q),$$

which is the multiplier robust control problem (3.4). We can drop the term $-\theta^* \eta$ from the objective without affecting the extremizing choices of (c, q) because we are holding θ fixed at θ^* .

Claim 3.6. *Suppose that for $\eta = \eta^*$, c^* and q^* solve the constraint robust control problem. Then there exists a $\theta^* \in \Theta$ such that the multiplier and constraint robust control problems have the same solution.*

Proof. This result is essentially the same as Theorem 2.1 of Petersen, James, and Dupuis (2000) and follows directly from Luenberger (1969). \square

Luenberger (1969) describes the following algorithm for constructing the multiplier. Let $J(c^*, \eta)$ satisfy:

$$J(c^*, \eta) = \inf_{q \in Q} \int_0^\infty \exp(-\delta t) \left(\int v_t(c^*) dq_t \right) dt$$

subject to $\mathcal{R}^*(q) \leq \eta$. As argued by Luenberger (1969), $J(c^*, \eta)$ is decreasing and convex in η . Given η^* , we let θ^* be the negative of the slope of the subgradient of $J(c^*, \cdot)$ at η^* . In other words, θ^* is the absolute value of the slope of a line tangent to $J(c^*, \cdot)$ at η^* .

This argument shows how to construct θ^* given η^* . It also suggests how to reverse the process. Given θ^* , we find a line with slope $-\theta^*$ that lies below $J(c^*, \cdot)$ and touches $J(c^*, \cdot)$ at a point η^* .

This argument fails however to account for the fact that the optimized choice of c may change as we alter η . Replacing $J(c^*, \cdot)$ by J^* from Definition 3.5 accounts for the optimization with respect to c and can sometimes be used in the construction. To study this further, consider the maximized objective \tilde{J} from Definition 3.4. Then

$$\tilde{J}(\theta) = \max_{c \in C} \min_{\eta \geq 0} J(c, \eta) + \theta \eta.$$

We shall study the consequences of the following assumption:

Assumption 3.7. *For any $\theta > \underline{\theta}$*

$$\begin{aligned} \tilde{J}(\theta) &= \max_{c \in C} \min_{q \in Q} \int_0^\infty \exp(-\delta t) \left(\int v_t(c) dq_t \right) dt + \theta \mathcal{R}^*(q) \\ &= \min_{q \in Q} \max_{c \in C} \int_0^\infty \exp(-\delta t) \left(\int v_t(c) dq_t \right) dt + \theta \mathcal{R}^*(q). \end{aligned}$$

Both equalities assume that the maximum and minimum are attained. More generally, we expect the second equality to be replaced by \leq because minimization occurs first. Section 5 tells how to verify Assumption 3.7 and discusses some of its ramifications.

Notice that

$$\begin{aligned}
\tilde{J}(\theta) &\doteq \max_{c \in \mathcal{C}} \min_{q \in \mathcal{Q}} \int_0^\infty \exp(-\delta t) \left(\int v_t(c) dq_t \right) dt + \theta \mathcal{R}^*(q) \\
&\leq \min_{\eta \geq 0} \max_{c \in \mathcal{C}} J(c, \eta) + \theta \eta \\
&\leq \min_{\eta \geq 0} J^*(\eta) + \theta \eta \\
&= \min_{q \in \mathcal{Q}} \max_{c \in \mathcal{C}} \int_0^\infty \exp(-\delta t) \left(\int v_t(c) dq_t \right) dt + \theta \mathcal{R}^*(q).
\end{aligned} \tag{8}$$

When Assumption 3.7 is satisfied, all of the inequalities become equalities, and

$$\tilde{J}(\theta) = \min_{\eta \geq 0} J^*(\eta) + \theta \eta. \tag{9}$$

Earlier we showed that the function $J(c^*, \eta)$ is convex and decreasing in η . Because it is the maximum of decreasing convex functions, the function J^* in Definition 3.5 is decreasing and convex in η . Equality (9) shows that \tilde{J} is the Legendre transform of J^* , which is known to be increasing and concave. The Legendre transform can be inverted to recover J^* from \tilde{J} :

$$J^*(\eta) = \max_{\theta \geq \underline{\theta}} \tilde{J}(\theta) - \eta \theta. \tag{10}$$

For a value of $\theta^* > \underline{\theta}$, formula (9) gives a corresponding value η^* as a solution to a maximization problem. For this η^* , formula (10) guarantees that there is a solution q^* with relative entropy η^* .

Claim 3.8. *Suppose that Assumption 3.7 is satisfied and that for $\theta > \underline{\theta}$, c^* is the maximizing choice of c for the multiplier robust control problem 3.4. Then that c^* also solves the constraint robust control problem 3.5 for $\eta = \eta^* = \mathcal{R}^*(q^*)$ where η^* solves (9).*

Claims 3.6 and 3.8 fully describe the mapping between the magnitudes of the constraint η and the multiplier θ . However, given η^* , they do not imply that the implied θ^* is unique, nor for a given $\theta^* > \underline{\theta}$ do they imply that the implied η^* is unique. While Claim 3.8 maintains Assumption 3.7, Claim 3.6 does not. Thus, without Assumption 3.7, for some values of θ a solution pair (c^*, q^*) of the multiplier problem cannot necessarily be interpreted as a solution to the constraint problem. Nevertheless, it suffices to limit attention to the family of multiplier problems because for any constraint, we can find a multiplier problem with the same solution pair (c^*, q^*) .

Section 5 describes sufficient conditions for Assumption 3.7. Our interest in these conditions extends beyond Claim 3.8 because they are informative about when solutions to the multiplier problem are recursive.

3.5 Bayesian Interpretation

Although Gilboa and Schmeidler (1989) suggest an alternative to the Bayesian notion of optimality, it is helpful to understand how a concern for robustness produces an implicit prior distribution that could be used to justify a robust decision rule in Bayesian terms. Assumption 3.7 allows us to construct a model for the shock process B under which the robust decision rule is optimal in a Bayesian sense.

Consider the control problem:

$$\max_{c \in \mathcal{C}} \int_0^\infty \exp(-\delta t) \left(\int v_t(c) dq_t \right) dt.$$

This problem takes as given the distortion q in the distribution of the Brownian motion. The optimal choice of a progressively measurable c responds to this distortion but does not presume to influence it. This optimized solution for c is not altered by adding $\theta \mathcal{R}^*(q)$ to the objective. Thus Assumption 3.7 allows us to support a solution to the multiplier problem by a particular distortion in the Brownian motion. The implied least favorable q^* is a valid probability distribution for the exogenous stochastic process $\{B_t : t \geq 0\}$, and c^* is the ordinary (non robust) optimal control process given that distribution. In the language of Bayesian decision theory, we can depict c^* as a Bayesian solution for a particular prior distribution over $\{B_t : t \geq 0\}$. (See Blackwell and Girshick (1954) and Chamberlain (2000) for related discussions.)

A similar argument applies to the constraint version of the robust control problem. Since the maximum of convex functions is convex,

$$\max_{c \in \mathcal{C}} \int_0^\infty \exp(-\delta t) \left(\int v_t(c) dq_t \right) dt + \theta \mathcal{R}^*(q)$$

is convex in q . From the Legendre transform,

$$J^*(\eta) = \max_{\theta \geq \underline{\theta}} \min_{q \in \mathcal{Q}} \max_{c \in \mathcal{C}} \int_0^\infty \exp(-\delta t) \left(\int v_t(c) dq_t \right) dt + \theta [\mathcal{R}^*(q) - \eta].$$

The parameter θ can now be interpreted as a Lagrange multiplier on the entropy constraint and is optimized to produce the worst-case distribution q^* that respects this constraint. Again we can view the optimized control process c^* from the inner-most maximization as a Bayesian solution to the control problem.

4 Parameterizing Perturbations

For conceptual and computational simplicity, we reformulate the multiplier version of a robust control problem as a two-player, zero sum, stochastic differential game. This allows us to choose a probability measure and a control sequence recursively. We accomplish this by reverting to the original probability space (Ω, \mathcal{F}, P) and by representing the alternative

probability specifications as either martingales or perturbations on this space. The martingale formulation compels us to add a state variable to the control problem posed in section 2 and a second player who influences this state variable. The closely related perturbation formulation avoids the use of an additional state variable, but omits some of the distributions that are contemplated by the multiplier problem.

4.1 Martingales

Recall that we consider only those probability distributions on $(\Omega^*, \mathcal{F}^*)$ that are absolutely continuous with respect to Wiener measure over finite intervals. Such absolute continuity is sufficient for us to obtain a martingale characterization of q . Let

$$\kappa_t(f) = \left(\frac{dq_t}{dq_t^0} \right) (f)$$

for any continuous function f in Ω^* . Here κ_t is the Radon-Nikodym derivative of q_t with respect to q_t^0 . Construct:

$$z_t = \kappa_t(B) \tag{11}$$

Claim 4.1. *Suppose that q is absolutely continuous with respect to q^0 . The process $\{z_t : t \geq 0\}$ defined via (11) on (Ω, \mathcal{F}, P) is a nonnegative martingale adapted to the filtration $\{\mathcal{F}_t : t \geq 0\}$ with $Ez_t = 1$. Moreover,*

$$\int \kappa_t \phi_t dq_t = E [z_t \phi_t(B)] \tag{12}$$

for any bounded and \mathcal{F}_t^* measurable function ϕ_t . Conversely if $\{z_t : t \geq 0\}$ is a nonnegative progressively measurable martingale with $Ez_t = 1$, then the probability measure q defined via (12) is absolutely continuous over finite intervals.

Proof. The first part of this claim follows directly from the proof of Theorem 7.5 in Lipster and Shiryaev (2000). This proof is essentially a direct application of the Law of Iterated Expectations and the fact that probability distributions necessarily integrate to one. Conversely, suppose that z is a nonnegative martingale on (Ω, \mathcal{F}, P) with unit expectation. Let ϕ_t be any nonnegative, bounded and \mathcal{F}_t^* measurable function. Define:

$$\int \phi_t dq_t = E [z_t \phi_t(B)].$$

This defines a measure because indicator functions are nonnegative, bounded functions. Clearly $\int \phi_t dq_t = 0$ whenever $E\phi_t(B) = 0$. Thus q_t is absolutely continuous with respect to q_t^0 , the measure induced by Brownian motion restricted to $[0, t]$. Setting $\phi_t = 1$ shows that q_t is in fact a probability measure for any t . \square

From this result we can perform integration on $(\Omega^*, \mathcal{F}^*, q)$ by integrating against a martingale z on the original probability space (Ω, \mathcal{F}, P) . In what follows, we will add a state variable z to the robust formulation of the control problem. To prepare the way, we must portray a martingale in a convenient way. Before proceeding, we note the following:

Remark 4.2. *A nonnegative martingale z with unit expectation can also be used to define a probability distribution on (Ω, \mathcal{F}) . For any random variable y_t that is \mathcal{F}_t measurable, we can build a probability measure on (Ω, \mathcal{F}_t) using z_t as a Radon-Nikodym derivative. For the results in section 3, it was most convenient to work with induced distributions. A martingale formulation for absolutely continuous probability distributions on the original measurable space (Ω, \mathcal{F}) is more convenient for solving the control problem and demonstrating recursivity.*

4.2 Representation

We attain a convenient representation of a martingale by exploiting the Brownian motion information structure. Any martingale M with a unit expectation can be portrayed as

$$z_t = 1 + \int_0^t k_u dB_u \quad (13)$$

where k is a progressively measurable d -dimensional process that satisfies:

$$P \left\{ \int_0^t |k_u|^2 du < \infty \right\} = 1$$

for any finite t (see Revuz and Yor (1994), Theorem V.3.4). Define:

$$h_t = \begin{cases} k_t/z_t & \text{if } z_t > 0 \\ 0 & \text{if } z_t = 0. \end{cases} \quad (14)$$

Then z solves the integral equation

$$z_t = 1 + \int_0^t z_u h_u dB_u$$

and its differential counterpart

$$dz_t = z_t h_t dB_t \quad (15)$$

with initial condition $z_0 = 1$ where

$$P \left\{ \int_0^t (z_u)^2 |h_u|^2 du < \infty \right\} = 1.$$

The scaling by $(z_u)^2$ permits

$$\int_0^t |h_u|^2 du = \infty$$

provided that $z_t = 0$ on this event.

We will use evolution equation (15) and index different martingales by different progressively measurable h 's. While a nonnegative martingale solves a stochastic differential equation (15) for an appropriate h , for some progressively measurable h 's for which (15) has a solution, the solution might be a supermartingale rather than a martingale. That is, the resulting z may only satisfy the inequality

$$E(z_t | \mathcal{F}_s) \leq z_s$$

for $0 < s \leq t$ even though z is a local martingale. For our robust control problems, we will need to verify that the associated solution for h implies that the resulting z is a martingale rather than the weaker supermartingale requirement.

4.3 Relative Likelihoods

We now study relative likelihoods under both the Brownian motion model and the alternative associated with a nonnegative martingale z . This is a necessary step in constructing our measure of relative entropy. First we consider the relative likelihood under the Brownian motion model for B .

The solution to (15) can be represented as an exponential:

$$z_t = \exp \left(\int_0^t h_u \cdot dB_u - \frac{1}{2} \int_0^t |h_u|^2 du \right). \quad (16)$$

But we must interpret this notation carefully. Since we allow for $\int_0^t |h_u|^2 du$ to be infinite with positive probability, we adopt the convention that the exponential is zero when this event happens. In the event that $\int_0^t |h_u|^2 du < \infty$, we can define the stochastic integral $\int_0^t h_u dB_u$ as an appropriate probability limit [see Lemma 6.2 of (Lipster and Shiryaev 2000)].

This exponential formula leads naturally to a relative likelihood. Thinking of B as being observed over an interval of time, we can construct a likelihood as a function of these data. Before approaching this construction, consider a simple motivating example:

Example 4.3. Consider two **models** of a vector y . In the first, y is normally distributed with mean ν and covariance matrix I . In the second y is normally distributed with mean zero and covariance matrix I . The logarithm of the ratio of first density to the second is:

$$\ell(y) = \left(\nu \cdot y - \frac{1}{2} \nu \cdot \nu \right).$$

Let E^1 denote the expectation under model one and E^2 two. Then

$$E^1 \exp[\ell(y)] = 1,$$

which follows from properties of the log-normal distribution. Under the second model

$$E^2 \ell(y) = E^1 \ell(y) \exp[\ell(y)] = \frac{1}{2} \nu \cdot \nu,$$

which is a measure of relative entropy.

When z is a martingale, we can interpret the left side of (16) as a formula for the relative likelihood of two models evaluated under the Brownian motion specification for B . Taking logarithms, we find that

$$\ell_t = \int_0^t h_u \cdot dB_u - \frac{1}{2} \int_0^t |h_u|^2 du.$$

Since h is progressively measurable, we can write:

$$h_t = \psi_t(B).$$

Changing the distribution of B in accordance with q results in another characterization of the relative likelihood. Let z be the martingale associated with q . The Girsanov Theorem implies that

$$\tilde{B}_t = B_t - \int_0^t h_u du \tag{17}$$

is a Brownian motion with respect to the filtration $\{\mathcal{F}_t : t \geq 0\}$ where h is defined by (14). Thus we can write the logarithm of the relative likelihood as:

$$\tilde{\ell}_t = \int_0^t \psi_u(B) \cdot d\tilde{B}_u + \frac{1}{2} \int_0^t |\psi_u(B)|^2 du. \tag{18}$$

The next claim asserts that both integrals on the right side are well-defined and finite, the first being a stochastic integral against a Brownian motion. This likelihood construction will allow us to obtain an alternative representation of relative entropy.

Claim 4.4. *Suppose that q_t is absolutely continuous with respect to q_t^0 for all $0 < t < \infty$. Let z be the corresponding nonnegative martingale on (Ω, \mathcal{F}, P) with $Ez_t = 1$. Then*

$$Ez_t \mathbf{1}_{\{\int_0^t |h_u|^2 du < \infty\}} = 1$$

where h_t is given by (14). Moreover,

$$\int \log \left(\frac{dq_t}{dq_t^0} \right) dq_t = \frac{1}{2} E \int_0^t z_u |h_u|^2 du.$$

Proof. See appendix A. □

The first result in Claim 4.4 guarantees that formula (18) is well defined, while the second result allows us to represent relative entropy of the induced distributions on the original probability space. Applying this claim, we can depict the entropy [formula (4)] of the measure q as:

$$\begin{aligned} \mathcal{R}^*(q) &= \delta \int_0^\infty \exp(-\delta t) \int \log \left(\frac{dq_t}{dq_t^0} \right) dq_t dt \\ &= \frac{\delta}{2} E \left[\int_0^\infty \exp(-\delta t) \int_0^t z_u |h_u|^2 du dt \right] \end{aligned}$$

$$= \frac{1}{2} E \left[\int_0^\infty \exp(-\delta t) z_t |h_t|^2 dt \right]$$

where the last equality follows from integrating by parts. This motivates our definition of entropy applied to nonnegative martingales:

$$\mathcal{R}(z) = \frac{1}{2} E \left[\int_0^\infty \exp(-\delta t) z_t |h_t|^2 dt \right] \quad (19)$$

This representation along with stochastic differential equation (15) allows us to formulate the robust control problem as a recursive stochastic differential game. Because we allow for nonnegative supermartingales that solve (15) for some h , we extend the entropy construction using (19) for some progressively measurable h .

Remark 4.5. Consider the process $\{z_t \log z_t : t \geq 0\}$, which is well defined until the process z hits zero. The drift for this process is $\frac{1}{2} z_t |h_t|^2$, the essential ingredient in our construction of relative entropy as we saw in (19).

Remark 4.6. If we had limited the process $h = \psi(B)$ so that the undiscounted integral

$$\int_0^\infty |\psi_t(B)|^2 dt < \infty \quad (20)$$

with probability one under q , then q would necessarily be absolutely continuous with respect to q^0 on $(\Omega^*, \mathcal{F}^*)$ [see Kabanov, Lipcer, and Sirjaev (1979)]. Any set that has Wiener measure zero, would also have q measure zero including sets defined in terms of the entire sample path of B . For instance, consider the large interval behavior of $\frac{1}{N} \int_0^N \phi_u(B) du$. If this obeys a Strong Law of Large Numbers under q^0 , then it will under q as well. Discounting directs the decision maker's attention away from tail events and toward the more immediate future, justifying a concern for misspecifications that are difficult to detect from finite amounts of data. For discounted problems, we study misspecifications that violate (20). We deliberately expand the set of misspecifications that concern the decision maker to include ones that can be difficult to detect from large but finite amounts of data.

4.4 A Penalized Martingale Problem

We now give an alternative specification of the multiplier version of the robust control problem, which we call a martingale robust control problem. Let H denote the set of d -dimensional progressively measurable processes.

Definition 4.7. A martingale robust control problem is:

$$\sup_{c \in C} \inf_{h \in H} E \left(\int_0^\infty \exp(-\delta t) z_t \left[U(c_t, x_t) + \frac{\theta}{2} |h_t|^2 \right] dt \right) \quad (21)$$

subject to:

$$\begin{aligned} dx_t &= \mu(c_t, x_t)dt + \sigma(c_t, x_t)dB_t \\ dz_t &= h_t z_t dB_t \end{aligned} \tag{22}$$

with initial conditions $x_0, z_0 = 1$.

To allow for alternative probability distributions we have added z_t as a multiplicative *preference shock* to the objective function. This process scales both the instantaneous utility and the quadratic penalty term $\frac{\theta}{2}|h_t|^2$.

We shall be interested in cases in which the following counterpart to Assumption 3.7 is satisfied.

Assumption 4.8.

$$\begin{aligned} \max_{c \in C} \min_{h \in H} E \left(\int_0^\infty \exp(-\delta t) z_t \left[U(c_t, x_t) + \frac{\theta}{2} |h_t|^2 \right] dt \right) = \\ \min_{h \in H} \max_{c \in C} E \left(\int_0^\infty \exp(-\delta t) z_t \left[U(c_t, x_t) + \frac{\theta}{2} |h_t|^2 \right] dt \right) \end{aligned}$$

subject to (22) with initial conditions $x_0, z_0 = 1$.

In the next section, we will describe how to verify this assumption.

Since the martingale control problem admits some supermartingales in addition to all nonnegative martingales for the process z , it is necessary that we verify that this enlargement fails to alter the solution of the control problem. In particular, consider the second problem in Assumption 4.8. Here it suffices to show that the minimizing h implies a z that is a martingale. Later we describe ways to check this condition.

4.5 A Penalized Perturbation Problem

Instead of changing the distribution via a martingale, suppose that we perturb the stochastic process used to model the shocks. That is, suppose that we replace the Brownian motion of B by the Ito perturbation

$$B_t = \int_0^t h_u du + \tilde{B}_t \tag{23}$$

where we assume that h is progressively measurable and \tilde{B} is a Brownian motion adapted to a filtration $\{\mathcal{F}_t : t \geq 0\}$. We make this specification on a probability space $(\Omega, \mathcal{F}, \tilde{P})$ and for the time being do not require that $\{\mathcal{F}_t : t \geq 0\}$ be the filtration generated by either \tilde{B} or B . Later we will explain the role of this added flexibility and restrict the filtration.

We limit the magnitude of the perturbation by using a quadratic penalty in h :

$$\frac{1}{2} \int_0^\infty \exp(-\delta t) \tilde{E} |h_t|^2 dt, \tag{24}$$

where $\tilde{E}(\cdot)$ denotes an expectation with respect to \tilde{P} . This gives rise to a second well-posed multiplier problem:

Definition 4.9. A *perturbed robust control problem* is:

$$\sup_{c \in C} \inf_{h \in H} \tilde{E} \left(\int_0^\infty \exp(-\delta t) \left[U(c_t, x_t) + \frac{\theta}{2} |h_t|^2 \right] dt \right) \quad (25)$$

subject to:

$$dx_t = \mu(c_t, x_t)dt + \sigma(c_t, x_t)(h_t dt + d\tilde{B}_t) \quad (26)$$

with initial condition x_0 .

This game alters the stochastic process used to model the shocks by appending a drift distortion that is disguised by the Brownian motion, leading to a stochastic version of a robust game analyzed by James (1992). James used the quadratic penalty in the drift distortion but studied a deterministic counterpart. The applications that interest us require an explicitly stochastic structure, like the ones that appear in the continuous-time formulation in Anderson, Hansen, and Sargent (2000) and the discrete-time formulation in Petersen, James, and Dupuis (2000).

The counterpart to Assumption 3.7 is:

Assumption 4.10.

$$\begin{aligned} \max_{c \in C} \min_{h \in H} \tilde{E} \left(\int_0^\infty \exp(-\delta t) \left[U(c_t, x_t) + \frac{\theta}{2} |h_t|^2 \right] dt \right) = \\ \min_{h \in H} \max_{c \in C} \tilde{E} \left(\int_0^\infty \exp(-\delta t) \left[U(c_t, x_t) + \frac{\theta}{2} |h_t|^2 \right] dt \right) \end{aligned}$$

subject to (26) with initial conditions x_0 .

The perturbed robust control problem is closely related to the martingale robust control problem. The martingale problem uses nonnegative martingales to parameterize alternative distributions on (Ω, \mathcal{F}) and hence alternative induced distributions on $(\Omega^*, \mathcal{F}^*)$. In the perturbed robust control problem, the Brownian motion B is replaced by a Brownian motion with a drift. This drift distortion induces an alternative distribution on $(\Omega^*, \mathcal{F}^*)$. The Girsanov Theorem used to justify (17) suggests that these two approaches to robustness are related, as we show next.

In the perturbed robust control problem, we limit perturbations using (24). When this discrepancy measure is finite, it follows that $\tilde{P}\{\int_0^t |h_u|^2 du < \infty\} = 1$. Theorem 7.4 of Lipster and Shiryaev (2000) implies that the distribution induced by B is absolutely continuous with respect to Wiener measure over finite intervals of time. Thus the Ito perturbation (23) gives a way to parameterize induced distributions that are absolutely continuous with respect to Wiener measure.

Two potential complications affect this parameterization. First, how do we construct the filtration used to restrict control processes and perturbations? Suppose the filtration is

generated by \tilde{B} . The process B built in (23) may generate a smaller filtration than \tilde{B} . When B and \tilde{B} happen to generate the filtration:

$$\delta \int_0^\infty \exp(-\delta t) \int \log \left(\frac{dq_t}{dq_t^0} \right) dq_t dt = \tilde{E} \int_0^\infty \exp(-\delta t) |h_t|^2 dt, \quad (27)$$

where q is the distribution induced by B . When B generates a smaller filtration, however, we can only say the right side of (27) is an upper bound for our measure of relative entropy. While this possibility is admitted in the problem formulation, we will be interested in solutions for which B generates $\{\mathcal{F}_t : t \geq 0\}$.

A second complication is that not all probability distributions that are absolutely continuous over finite intervals can be constructed using Ito perturbations of the form (23). The Girsanov Theorem implies the following weaker result, which gives the perturbation counterpart to Claim 4.4:

Claim 4.11. *If q is absolutely continuous with respect to q^0 , then q is the induced distribution for a (possibly weak) solution B to a stochastic differential equation defined on a probability space $(\Omega, \mathcal{F}, \tilde{P})$:*

$$dB_t = \psi_t(B)dt + d\tilde{B}_t \quad (28)$$

for some progressively measurable ψ defined on $(\Omega^*, \mathcal{F}^*)$ and some Brownian motion \tilde{B} that is adapted to $\{\mathcal{F}_t : t \geq 0\}$. Moreover, for each t

$$\tilde{P} \left[\int_0^t |\psi_u(B)|^2 du < \infty \right] = 1.$$

Proof. From Lemma 4.1 there is a nonnegative martingale z associated with the Radon-Nikodym derivative of q with respect to q^0 . This martingale has expectation unity for all t . The conclusion follows from a generalization of the Girsanov Theorem (*e.g.* see Lipster and Shiryaev (2000) Theorem 6.2). \square

Since the solution to stochastic differential equation (28) is not necessarily a strong solution, B may generate a *larger* filtration than \tilde{B} .⁶ This leads us to consider weak solutions to the penalized perturbation problem.

Remark 4.12. *The conclusion that for each t ,*

$$\tilde{P} \left[\int_0^t |\psi_u(B)|^2 du < \infty \right] = 1$$

in Claim 4.11 is the counterpart to the conclusion:

$$E z_t \mathbf{1}_{\{\int_0^t |h_u|^2 du < \infty\}} = 1$$

⁶See Tsirel'son (1975) for an example of a bounded ψ for which (28) fails to have a strong solution. In his example, B generates a strictly larger filtration than \tilde{B} although \tilde{B} remains a Brownian motion with respect to this larger filtration. The induced distributions that emerge from his example are equivalent (mutually absolutely continuous).

for $t > 0$ in Claim 4.4. Consider a stochastic integral:

$$\int_0^t \phi_u(B) \cdot dB_u$$

which is well defined on (Ω, \mathcal{F}, P) provided that $\int_0^t |\phi_u(B)|^2 du < \infty$ with P probability one. Then

$$\int_0^t \phi_u(B) \cdot dB_u = \int_0^t \psi_u(B) \cdot \phi_u(B) du + \int_0^t \phi_u(B) \cdot d\tilde{B}_u$$

is well defined on $(\Omega, \mathcal{F}, \tilde{P})$ because $\int_0^t |\phi_u(B)|^2 du < \infty$ with \tilde{P} probability one (absolute continuity) and $\int_0^t |\psi_u(B)|^2 du < \infty$ with \tilde{P} probability one as implied by Claim 4.4.

An advantage of the perturbation formulation over the martingale formulation is that we avoid having to introduce a new state variable. In what follows we will use the solution to the martingale problem to construct a solution to the perturbation problem in which the filtration used to restrict decisions is generated by B . In particular, we will assume that we have strong solution to the martingale problem but possibly a weak solution to the perturbation problem. Because the Girsanov Theorem only justifies a weak solution, for the perturbation problem we will allow the Brownian motion to generate a smaller information set than the Markov state.

5 Recursivity of Multiplier Formulations

Building on a result from Fleming and Souganidis (1989), this section studies the recursivity of the martingale robust control problem and thereby establishes a direct way to show the recursivity of the multiplier robust control problem. The next section then shows how the connections between the multiplier and constraint control problems make the recursivity of the multiplier problem carry over to the constraint problem.

The robust martingale problem 4.7 is a special case of the two-player, zero-sum, stochastic differential games studied by Fleming and Souganidis (1989). The martingale problem 4.7 assumes that at time zero both decision makers commit to decision processes whose time t components are measurable functions of \mathcal{F}_t . The decision maker who chooses distorted beliefs $\{h_t\}$ takes $\{c_t\}$ as given; and the decision maker who chooses $\{c_t\}$ takes $\{h_t\}$ as given. Assumption 3.7 asserts that the order in which the two decision makers choose these processes does not matter: the date zero value function is unaffected by which decision maker chooses first.

This description requires that at time zero both decision makers commit to their respective decision processes. We now alter the timing protocols and explore conditions under which allowing the two players to choose sequentially implies the same time zero value function for the game. As a by-product, our argument will justify the exchange of orders of extremization stipulated by Assumption 3.7.

We have used c to denote the control process. We now use \check{c} to denote the value of a control at a particular date. In the recursive formulation, we restrict \check{c} to be in some set \check{C} common for all dates. This imposes more structure than we have before on the set C of admissible control processes. We let \check{h} denote the realized martingale control at any particular date. We can think of \check{h} as a vector in \mathbb{R}^d . Similarly, we think of \check{x} and \check{z} as being realized states.

We now let the initial state vary and define a value function \tilde{V} as the objective for the martingale problem. In particular, we shall verify that $\tilde{J}(\theta) = \check{z}\tilde{V}(\check{x}, \theta)$, provided that \check{x} is initialized at x_0 and \check{z} is initialized at one. With a recursive solution, this same value function is valid in subsequent time periods. The linearity of J in \check{z} must be verified.

Fleming and Souganidis (1989) study when there exists a recursive solution to the multiplier problem 4.7.⁷ They use a Bellman-Isaacs condition to justify a recursive solution, that is, to render equilibrium outcomes under two-sided commitment at date zero identical with outcomes of a Markov perfect equilibrium in which the decision rules of both agents are recursively chosen to be functions of the state vector x_t .

The Hamilton-Jacobi-Bellman equation for a Markov counterpart to the martingale game is:

$$\begin{aligned} \delta\check{z}\tilde{V}(\check{x}, \theta) &= \max_{\check{c} \in \check{C}} \min_h \check{z}U(\check{c}, \check{x}) + \check{z}\frac{\theta}{2}\check{h} \cdot \check{h} + \mu(\check{c}, \check{x}) \cdot \tilde{V}_x(\check{x}, \theta)\check{z} \\ &\quad + \check{z}\frac{1}{2}\text{trace} \left[\sigma(\check{c}, \check{x})'\tilde{V}_{xx}(\check{x}, \theta)\sigma(\check{c}, \check{x}) \right] + \check{z}\check{h} \cdot \sigma(\check{c}, \check{x})'\tilde{V}_x(\check{x}, \theta) \end{aligned} \quad (29)$$

where \tilde{V}_x is the vector of partial derivatives of \tilde{V} with respect to \check{x} and \tilde{V}_{xx} is the matrix of second derivatives. Equation (29) is a Bellman equation for an infinite-horizon zero-sum two-player game. The diffusion specification makes this Bellman equation a partial differential equation. It has multiple solutions that correspond to different boundary conditions. To find the actual value function requires that we apply a *Verification Theorem* (e.g. see Theorem 5.1 of Fleming and Soner (1993)).

The partial differential equation (29) is scaled by \check{z} , which verifies our guess that the

⁷Fleming and Souganidis (1989) impose as restrictions that μ , σ and U are bounded, uniformly continuous and Lipschitz continuous with respect to \check{x} uniformly in \check{c} . They also require that the controls \check{c} and \check{h} reside in compact sets. While these restrictions are imposed to obtain general existence results, they are not satisfied for some important examples. Presumably existence in these examples will require special arguments. These issues are beyond the scope of this paper.

Furthermore, it is known that in general the value functions associated with stochastic control problems will not be twice differentiable, as would be required for the Hamilton-Jacobi-Bellman equations in Assumption 5.1 below to possess classical solutions. However Fleming and Souganidis (1989) prove that the value function satisfies the Hamilton-Jacobi-Bellman equation in a weaker *viscosity* sense. Viscosity solutions are often needed when it is feasible and sometimes desirable to set the control \check{c} so that $\sigma(\check{c}, \check{x})$ has lower rank than d , which is the dimension of the Brownian motion.

value function is linear in z and allows us to study the alternative equation:

$$\begin{aligned} \delta\tilde{V}(\check{x}, \theta) &= \max_{\check{c} \in \check{C}} \min_{\check{h}} U(\check{c}, \check{x}) + \frac{\theta}{2} \check{h} \cdot \check{h} + [\mu(\check{c}, \check{x}) + \sigma(\check{c}, \check{x})\check{h}] \cdot \tilde{V}_x(\check{x}, \theta) \\ &\quad + \frac{1}{2} \text{trace} \left[\sigma(\check{c}, \check{x})' \tilde{V}_{xx}(\check{x}, \theta) \sigma(\check{c}, \check{x}) \right] \end{aligned} \quad (30)$$

This partial differential equation only involves the state vector \check{x} and not \check{z} . Note also that \tilde{V} is the value function for the recursive version of the perturbed robust control game.

The Bellman-Isaacs condition for this differential game is:

Assumption 5.1. *The value function \tilde{V} satisfies*

$$\begin{aligned} \delta\tilde{V}(\check{x}, \theta) &= \max_{\check{c} \in \check{C}} \min_{\check{h}} U(\check{c}, \check{x}) + \frac{\theta}{2} \check{h} \cdot \check{h} + [\mu(\check{c}, \check{x}) + \sigma(\check{c}, \check{x})\check{h}] \cdot \tilde{V}_x(\check{x}, \theta) \\ &\quad + \frac{1}{2} \text{trace} \left[\sigma(\check{c}, \check{x})' \tilde{V}_{xx}(\check{x}, \theta) \sigma(\check{c}, \check{x}) \right] \\ &= \min_{\check{h}} \max_{\check{c} \in \check{C}} U(\check{c}, \check{x}) + \frac{\theta}{2} \check{h} \cdot \check{h} + [\mu(\check{c}, \check{x}) + \sigma(\check{c}, \check{x})\check{h}] \cdot \tilde{V}_x(\check{x}, \theta) \\ &\quad + \frac{1}{2} \text{trace} \left[\sigma(\check{c}, \check{x})' \tilde{V}_{xx}(\check{x}, \theta) \sigma(\check{c}, \check{x}) \right] \end{aligned}$$

Fleming and Souganidis (1989) show that the freedom to exchange orders of maximization and minimization guarantees that equilibria of the date zero commitment and the Markov perfect multiplier games coincide. The ability to exchange orders of extremization in the recursive specification implies that the orders of extremization can also be exchanged in the date zero commitment problem, as required in Assumption 3.7. As we shall now see, the exchange of order of extremization in Assumption 5.1 can often be verified without precise knowledge of the value function \tilde{V} .

5.1 No Binding Inequality Restrictions

Suppose that there are no binding inequality restrictions on c . Then a justification for Assumption 5.1 can emerge from the first-order conditions for \check{c} and \check{h} . Define

$$\begin{aligned} \chi(\check{c}, \check{h}, \check{x}) &\doteq U(\check{c}, \check{x}) + \frac{\theta}{2} \check{h} \cdot \check{h} + [\mu(\check{c}, \check{x}) + \sigma(\check{c}, \check{x})\check{h}] \cdot \tilde{V}_x(\check{x}, \theta) \\ &\quad + \frac{1}{2} \text{trace} \left[\sigma(\check{c}, \check{x})' \tilde{V}_{xx}(\check{x}, \theta) \sigma(\check{c}, \check{x}) \right], \end{aligned} \quad (31)$$

and suppose that χ is continuously differentiable in \check{c} . First, find a Nash equilibrium by solving:

$$\begin{aligned} \frac{\partial \chi}{\partial c}(\check{c}^*, \check{h}^*, \check{x}) &= 0 \\ \frac{\partial \chi}{\partial h}(\check{c}^*, \check{h}^*, \check{x}) &= 0. \end{aligned}$$

In particular, the first-order conditions for \check{h} are:

$$\frac{\partial \chi}{\partial h}(\check{c}^*, \check{h}^*, \check{x}) = \theta \check{h}^* + \sigma(\check{c}^*, \check{x})' \tilde{V}_x(\check{x}, \theta) = 0.$$

If a unique solution exists and if it suffices for extremization, the Bellman-Isaacs condition is satisfied. This follows from the ‘‘chain rule.’’ Suppose that the minimizing player goes first and computes \check{h} as a function of \check{x} and \check{c} :

$$\check{h}^* = -\frac{1}{\theta} \sigma(\check{c}, \check{x})' \tilde{V}_x(\check{x}, \theta) \quad (32)$$

Then the first-order conditions for the max player selecting \check{c} as a function of \check{x} are:

$$\frac{\partial \chi}{\partial c} + \frac{\partial h}{\partial c} \frac{\partial \chi}{\partial h} = 0$$

where $\frac{\partial h}{\partial c}$ can be computed from the reaction function (32). Notice that the first-order conditions for the maximizing player are satisfied at the Nash equilibrium. A similar argument can be made if the maximizing player chooses first.

5.2 Separability

Consider next the case in which σ does not depend on the control. In this case the decision problems for \check{c} and \check{h} separate. For instance, from (32), we see that \check{h} does not react to \check{c} in the minimization of \check{h} conditioned on \check{c} . Even with binding constraints on \check{c} , the Bellman-Isaacs condition (Assumption 5.1) is satisfied, provided that a solution exists for \check{c} .

5.3 Convexity

A third approach that uses results of Fan (1952) and Fan (1953) is based on the global shape properties of the objective. When we can reduce the choice set C to be a compact subset of a linear space, Fan (1952) can apply. Fan (1952) also requires that the set of conditional minimizers and maximizers be convex. We know from formula (32) that the minimizers of $\chi(\check{c}, \cdot, \check{x})$ form a singleton set, which is convex for each \check{c} and \check{x} .⁸ Suppose also that the set of maximizers of $\chi(\cdot, \check{h}, \check{x})$ is non-empty and convex for each \check{h} and \check{x} .⁹ Then again the Bellman-Isaacs condition (Assumption 5.1) is satisfied. Finally Fan (1953) does not require that the set \check{C} be a subset of a linear space, but instead requires that $\chi(\cdot, \check{h}, \check{x})$ be concave. By relaxing the linear space structure we can achieve compactness by adding points (say the point ∞) to the control set, provided that we can extend $\chi(\cdot, \check{h}, \check{x})$ to be upper semi-continuous. The extended control space must be a compact Hausdorff space. Provided that

⁸Notice that provided \check{C} is compact, we can use (32) to specify a compact set that contains the entire family of minimizers for each \check{c} in \check{C} and a given \check{x} .

⁹See Ekeland and Turnbull (1983) for a discussion of continuous time, deterministic control problems when the set of minimizers is not convex. They show that sometimes it is optimal to *chatter* between different controls as a way to imitate convexification in continuous time.

the additional points are not attained in optimization, we can apply Fan (1953) to verify Assumption 5.1.¹⁰

6 Recursive Representation of the Commitment Equilibrium

We have discussed the connection between a robust control problem and a two-person zero-sum game in which at date zero both players commit to entire decision processes.¹¹ In that time-zero game, the decision maker's (*i.e.*, the maximizing player's) actions do not alter the distribution q . This was necessary for us to be able to represent the control problem as a Bayesian solution for a particular prior distribution. The Markov perfect equilibrium characterized by (29) has a different timing protocol that makes it less evident that the decision maker's actions don't influence the distribution q because the control law for h can depend on states that can be influenced by the control c . However, there is a way of interpreting the (constrained) worst-case model as one in which $\{B_t : t \geq 0\}$ cannot be influenced by the decision maker's choice of control. In particular, in this section we show how by using a version of the 'big X , little x ' trick common in macroeconomics, we can use the Markov solution to the robust multiplier game to get a recursive representation of the equilibrium under two-sided commitment at time-0 where the minimizing player first chooses q . The 'big X , little x ' trick allows us to derive an exogenous specification of $\{B_t : t \geq 0\}$ for which the control process from the robust perturbation game is optimal. We find a recursive version of the commitment solution by revisiting our discussion at the end of section 3.

Since we allowed for a supermartingale as a solution for the martingale problem, we also suggest a way to verify that the supermartingale is actually a martingale. This is needed for us to argue that the h solution to the martingale problem implies a well defined alternative probability model.

6.1 State Evolution

Suppose that a progressively measurable process c is chosen optimally given a process h that is used to model the martingale process z .¹² Under this view, the control c cannot influence future values of h . The determination of the constrained worst-case process $\{h_t : t \geq 0\}$ will depend on the initial state x_0 , but the evolution for the forcing process B cannot be influenced

¹⁰Apply Theorem 2 of Fan (1953) to $-\chi(\cdot, \cdot, \check{x})$. This theorem does not require compactness of choice set for \check{h} , only of the choice set for \check{c} . The theorem also does not require attainment when optimization is over the noncompact choice set. In our application, we can verify attainment directly.

¹¹See Hansen and Sargent (2001a) for a related discussion cast in terms of a discrete time version of a permanent income model.

¹²Under this timing protocol, Fleming and Souganidis (1989) refer to a decision rule making $\{c_t : t \geq 0\}$ depend on $\{h_t : t \geq 0\}$ as a strategy. In their language, a strategy maps one progressively measurable process into another one.

by c . The distorted evolution facing the decision-maker can be portrayed recursively as:¹³

$$\begin{aligned} dx_t &= \mu(c_t, x_t)dt + \sigma(c_t, x_t)dB_t \\ dz_t &= z_t\alpha_h(X_t)dB_t \\ dX_t &= \mu^*(X_t)dt + \sigma^*(X_t)dB_t. \end{aligned} \tag{33}$$

Here the decision maker views x as a potentially controllable part of the state and regards X and hence z as an uncontrollable part. In (33), $c_t = \alpha_c(x_t)$ and $h_t = \alpha_h(x_t)$ is the Markov solution of the differential game (30). The coefficients for the evolution of X satisfy

$$\begin{aligned} \mu^*(\check{X}) &= \mu[\alpha_c(\check{X}), \check{X}] \\ \sigma^*(\check{X}) &= \sigma[\alpha_c(\check{X}), \check{X}]. \end{aligned}$$

Thus the evolution of (X, z) agrees with the solution to the Markov martingale game.

By construction, the control process c is not allowed to influence the state vector process X in (33) and the drift distortion $h_t = \alpha_h(X_t)$ for the martingale z depends only on this uncontrollable state vector. The solution $c_t = \xi_c(x_t, X_t)$ to the Markov control problem will satisfy $\alpha_c(\check{x}) = \xi_c(\check{x}, \check{x})$. Provided that X_0 is initialized at the same value as x_0 , at the optimized solution $X_t = x_t$ because the implied stochastic evolution for the two state vector processes coincide.¹⁴

6.2 Martingale Solution

While the solution process z is a local martingale and a supermartingale, for it to correspond to an alternative probability assessment we require that z be a martingale. This is important for two reasons. In formulating the martingale problem we added in some local martingales into the choice set for the minimizing agent for technical convenience. We must check that this augmentation does not alter the solution. Since the Bellman-Isaacs condition 4.8 is satisfied, it suffices to check this for the decision order

$$\min_{h \in H} \max_{c \in C} E \left(\int_0^\infty \exp(-\delta t) z_t \left[U(c_t, x_t) + \frac{\theta}{2} |h_t|^2 \right] dt \right)$$

subject to (22) with initial conditions $x_0, z_0 = 1$.¹⁵ We also need the z solution to be a martingale to support a Bayesian interpretation.

¹³See the appendix B for more details about the ‘big X , little x ’ evolution equation (33) and the associated value function.

¹⁴This interpretation of (29) isolates the distortions from influence by the decision maker. We would not always want to insist on this interpretation. Indeed, as a vehicle for promoting robustness to misspecification of endogenous dynamics, there are contexts in which we may want to allow the decision maker to imagine that his choice of c affects future distortions h .

¹⁵To see this let $H^* \subseteq H$ be the set of controls h for which z is a martingale and let $\text{obj}(h, c)$ be the objective as a function of the controls. Then under a Bellman-Isaacs condition we have

$$\min_{h \in H^*} \max_{c \in C} \text{obj}(h, c) \geq \min_{h \in H} \max_{c \in C} \text{obj}(h, c) = \max_{c \in C} \min_{h \in H} \text{obj}(h, c) \leq \max_{c \in C} \min_{h \in H^*} \max_{c \in C} \text{obj}(h, c). \tag{34}$$

6.3 Interpretation

Under the perturbation interpretation, the formula for α_h

$$\alpha_h(\check{x}) = -\frac{1}{\theta} \sigma^*(\check{x})' \tilde{V}_x(\check{x}, \theta).$$

introduces a drift distortion that is directly related to how the value function stochastic process responds to a Brownian motion shock. A shock that shifts this value function positively is offset by a negative drift distortion. Similarly, we can get a formula for the martingale used in constructing the distorted probability distribution. Suppose that $\tilde{V}(\check{x}, \theta)$ is twice continuously differentiable. Applying the formula on page 226 of Revuz and Yor (1994), form the positive function:

$$W(\check{x}, \theta) = \exp \left[-\frac{1}{\theta} \tilde{V}(\check{x}, \theta) \right]$$

Then

$$z_t = \frac{W(X_t, \theta)}{W(X_0, \theta)} \exp \left[-\int_0^t w(X_u) du \right]$$

where w is used to ensure that z has a zero drift. The worst case distribution assigns more weight to bad states as measured by an exponential adjustment to the value function.

7 Risk Sensitivity

The Hamilton-Jacobi-Bellman equation (30) also arises from a risk-sensitive control problem. Risk sensitive optimal control was initiated by Jacobson (1973) and Whittle (1981) in the context of discrete-time linear-quadratic decision problems. Letting ρ be an intertemporal return function, instead of maximizing $E\rho$ (where E continues to mean mean mathematical expectation), risk-sensitive control theory maximizes $E[\exp(\theta^{-1}\rho)]$, where θ^{-1} is a risk-sensitivity parameter. Jacobson and Whittle showed that the risk-sensitive control law can be computed by solving a robust multiplier problem of the type we have studied here. Hansen and Sargent (1995) showed how to use recursive utility theory to introduce discounting into the linear-quadratic, Gaussian risk-sensitive decision problem. James (1992) studied a continuous-time, nonlinear diffusion formulation of a risk-sensitive control problem and its robust counterpart in the absence of discounting. Again, the control law that solves the risk-sensitive problem also solves a stochastic robust multiplier problem. As emphasized by Anderson, Hansen, and

If we demonstrate, the first inequality \geq in (34) is an equality, it follows that

$$\min_{h \in H^*} \max_{c \in C} \text{obj}(h, c) \leq \min_{h \in H^*} \max_{c \in C} \text{obj}(h, c).$$

Since the reverse inequality is always satisfied provided that the extrema are attained, this inequality can be replaced by an equality. It follows that the second inequality \leq in (34) must in fact be an equality as well.

Sargent (2000), this equivalence carries over to problems with discounting when a recursive counterpart to the risk-sensitive objective is used.

The Bellman equation for the recursive, risk-sensitive control problem is obtained by substituting the solution (32) for h into the partial differential equation (30):

$$\begin{aligned}
\delta\tilde{V}(\check{x}, \theta) &= \max_{\check{c} \in \check{C}} \min_{\check{h}} U(\check{c}, \check{x}) + \frac{\theta}{2} \check{h} \cdot \check{h} + [\mu(\check{c}, \check{x}) + \sigma(\check{c}, \check{x})\check{h}] \cdot \tilde{V}_x(\check{x}, \theta) \\
&\quad + \frac{1}{2} \text{trace} \left[\sigma(\check{c}, \check{x})' \tilde{V}_{xx}(\check{x}, \theta) \sigma(\check{c}, \check{x}) \right] \\
&= \max_{\check{c} \in \check{C}} U(\check{c}, \check{x}) + \mu(\check{c}, \check{x}) \cdot \tilde{V}_x(\check{x}, \theta) \\
&\quad + \frac{1}{2} \text{trace} \left[\sigma(\check{c}, \check{x})' \tilde{V}_{xx}(\check{x}, \theta) \sigma(\check{c}, \check{x}) \right] \\
&\quad - \frac{1}{2\theta} \tilde{V}_x(\check{x}, \theta)' \sigma(\check{c}, \check{x}) \sigma(\check{c}, \check{x})' \tilde{V}_x(\check{x}, \theta)
\end{aligned} \tag{35}$$

The term

$$\mu(\check{c}, \check{x}) \cdot \tilde{V}_x(\check{x}, \theta) + \frac{1}{2} \text{trace} \left[\sigma(\check{c}, \check{x})' \tilde{V}_{xx}(\check{x}, \theta) \sigma(\check{c}, \check{x}) \right]$$

in Bellman equation (35) is the local mean or dt contribution to the value function process $\{\tilde{V}(x_t, \theta) : t \geq 0\}$ without any reference to model misspecification. Thus (35) coincides with the Bellman equation for the benchmark control problem (2), (3), but with an additional term included:

$$-\frac{1}{2\theta} \tilde{V}_x(\check{x}, \theta)' \sigma(\check{c}, \check{x}) \sigma(\check{c}, \check{x})' \tilde{V}_x(\check{x}, \theta).$$

This term is familiar from the analysis of continuous-time, stochastic specifications of recursive utility by Duffie and Epstein (1992). Notice that $\tilde{V}_x(x_t, \theta)' \sigma(c_t, x_t) dB_t$ gives the local Brownian contribution to the value function process $\{\tilde{V}(x_t, \theta) : t \geq 0\}$. The additional term in the Bellman equation is the negative of the local variance of the continuation value weighted by $\frac{1}{2\theta}$. Thus the *risk sensitive* interpretation excludes worries about misspecified dynamics and instead enhances the control objective with aversion to risk in a way captured by the local variance of the continuation value.

Duffie and Epstein (1992) refer to $\frac{1}{\theta}$ as the variance multiplier. Notice that here the variance multiplier is state independent, which emerges because an exponential risk adjustment is made to the continuation value.¹⁶ As a consequence of this adjustment, the Bellman equation contains a contribution from the local variance of the continuation value function. Solving the Bellman equation for the robust multiplier problem is equivalent to solving the Bellman equation for the risk-sensitive problem. While mathematically similar to the situation discussed in James (1992) (see pages 403 and 404), the presence of discounting in our setup compels us to use a recursive representation of the objective of the decision-maker. In the remainder of this paper we will be primarily concerned with the robustness interpretation, although we will revisit the recursive formulation of risk-sensitivity when we discuss preference orders.

¹⁶This is analogous to the exponential risk adjustment used elsewhere in the risk-sensitive control literature.

8 Recursivity of the Constraint Formulation

This section shows that after we add an additional state variable and an additional vector of controls, the constraint robust control problem also has a recursive structure.

For the date zero constraint problem, we studied how the objective depended on the magnitude of the entropy constraint. Now at each date we must carry along a state variable that measures the entropy that remains to be allocated. Instead of a value function \tilde{V} that depends only on the state x , we now use a value function V^* that depends also on that additional state variable, denoted r .

8.1 An Alternative Bellman Equation

Our strategy will be to use (9) to link the value functions for the multiplier and constraint problems, then to deduce from the Bellman equation (30) a partial differential equation that can be interpreted as the Bellman equation for another two-player game with additional states and controls. By construction, that new game is recursive and will have the same equilibrium outcome and representation as game (30). We shall then interpret this new game as recursively solving our original robust constraint problem 3.5.

In section 3 we argued that the date zero value functions for the constraint and multiplier problems are related via the Legendre transform. This leads us to construct:

$$\tilde{z}V^*(\tilde{x}, \tilde{r}) = \max_{\theta \geq \tilde{\theta}} \tilde{z}\tilde{V}(\tilde{x}, \theta) - \tilde{r}\tilde{z}\theta \quad (36)$$

where $J^*(\eta) = V^*(\tilde{x}, \tilde{r})$ provided that \tilde{x} is equal to the date zero state x_0 , \tilde{r} is set to the initial entropy constraint η and $\tilde{z} = 1$. The Legendre transform (36) scales linearly in the nonnegative variable \hat{z} . In next subsections we will have cause to introduce an additional state variable r that we interpret as the continuation value of entropy. Before doing that, we deduce a partial differential equation for V^* . While $\tilde{z}V^*(\tilde{x}, \tilde{r})$ will be the value function for the martingale version, as with the multiplier game we may interpret V^* as the value function for a perturbation formulation.

Inverting (36) yields

$$\tilde{V}(\tilde{x}, \theta) = \inf_{\tilde{r} \geq 0} V^*(\tilde{x}, \tilde{r}) + \theta\tilde{r}.$$

The function V^* is convex in r and has $-\theta$ as a subgradient with respect to r . In particular, when V^* is differentiable in r ,

$$\frac{\partial V^*}{\partial r}(\tilde{x}, \tilde{r}) = -\theta. \quad (37)$$

Inverting this function gives r as a function of x for a given θ . By the Implicit Function Theorem,

$$\frac{\partial r}{\partial x} = -\frac{V_{rx}^*}{V_{rr}^*}.$$

Thus

$$\tilde{V}_x = V_x^* \quad (38)$$

and

$$\begin{aligned}\tilde{V}_{xx} &= V_{xx}^* + V_{xr}^* \frac{\partial r}{\partial x} \\ &= V_{xx}^* - \frac{V_{xr}^* V_{rx}^*}{V_{rr}^*}.\end{aligned}\tag{39}$$

The Hamilton-Jacobi-Bellman partial differential equation (30) for \tilde{V} implies a corresponding partial differential equation for V^* that can be deduced by using formulas (38) and (39). Notice first that

$$\begin{aligned}&\frac{1}{2}\text{trace} \left[\sigma(\check{c}, \check{x})' \tilde{V}_{xx}(\check{x}, \theta) \sigma(\check{c}, \check{x}) \right] = \\ \min_{\check{g}} &\frac{1}{2}\text{trace} \left(\begin{bmatrix} \sigma(\check{c}, \check{x})' & \check{g} \end{bmatrix} \begin{bmatrix} V_{xx}^*(\check{x}, \check{r}) & V_{xr}^*(\check{x}, \check{r}) \\ V_{rx}^*(\check{x}, \check{r}) & V_{rr}^*(\check{x}, \check{r}) \end{bmatrix} \begin{bmatrix} \sigma(\check{c}, \check{x}) \\ \check{g}' \end{bmatrix} \right).\end{aligned}$$

This equality follows because the solution to the minimization problem is

$$\check{g}^* = -\frac{\sigma(\check{c}, \check{x})' V_{xr}^*(\check{x}, \check{r})}{V_{rr}^*(\check{x}, \check{r})}$$

and (39) is satisfied. Thus game (30) implies that

$$\begin{aligned}\delta V^*(\check{x}, \check{r}) &= \max_{\check{c} \in \check{C}} \min_{\check{h}, \check{g}} U(\check{c}, \check{x}) + [\mu(\check{c}, \check{x}) + \sigma(\check{c}, \check{x})\check{h}] \cdot V_x^*(\check{x}, \check{r}) + \left(\delta \check{r} - \frac{\check{h} \cdot \check{h}}{2} \right) \cdot V_r^*(\check{x}, \check{r}) \\ &\quad + \frac{1}{2}\text{trace} \left(\begin{bmatrix} \sigma(\check{c}, \check{x})' & \check{g} \end{bmatrix} \begin{bmatrix} V_{xx}^*(\check{x}, \check{r}) & V_{xr}^*(\check{x}, \check{r}) \\ V_{rx}^*(\check{x}, \check{r}) & V_{rr}^*(\check{x}, \check{r}) \end{bmatrix} \begin{bmatrix} \sigma(\check{c}, \check{x}) \\ \check{g}' \end{bmatrix} \right),\end{aligned}\tag{40}$$

given (37).

Equation (40) supports a recursive formulation of the constraint game. In particular, (40) is interpretable as a Hamilton-Jacobi-Bellman equation with a new control g and a new state r with evolution:

$$\begin{aligned}dr_t &= \left(\delta r_t - \frac{h_t \cdot h_t}{2} - g_t \cdot h_t \right) dt + g_t \cdot dB_t \\ &= \left(\delta r_t - \frac{h_t \cdot h_t}{2} \right) dt + g_t \cdot d\tilde{B}_t.\end{aligned}\tag{41}$$

The first equation is used for a martingale specification of the probability distortion captured by a nonnegative martingale z . Recall that in the martingale problem B is a Brownian motion. The second equation is pertinent for the perturbation specification in which \tilde{B} is a Brownian motion and h is the drift distortion. Notice that V_r^* in the partial differential equation (40) multiplies the drift in the second evolution equation.

The control g_t in (41) is chosen by the minimizing agent. This interpretation is valid provided that we can show that (40) is satisfied along the solution trajectory for the implied game. Before addressing this point, the next section shows that (41) describes the evolution of the continuation value of relative entropy.

8.2 Recursivity of Relative Entropy

Building a recursive specification of the constraint problem requires a recursive depiction of relative entropy. For a given probability distribution q , let z be the corresponding martingale with date t drift $z_t h_t$. Recall formula (19) for relative entropy of a nonnegative martingale:

$$\mathcal{R}(z) \doteq E \int_0^\infty \exp(-\delta t) z_t |h_t|^2 dt.$$

We define a date t conditional counterpart as follows:

$$\mathcal{R}_t(z) \doteq E \left[\int_0^\infty \exp(-\delta u) \left(\frac{z_{t+u}}{z_t} \right) |h_{t+u}|^2 du \middle| \mathcal{F}_t \right]$$

provided that $z_t > 0$ and $\mathcal{R}_t(z)$ to be zero otherwise. Notice that $\mathcal{R}_t(z)$ is a random variable defined on the original probability space indexed by t . This family of random variables induces the recursion for $\epsilon > 0$:

$$z_t \mathcal{R}_t(z) = \exp(-\delta \epsilon) E \left[z_{t+\epsilon} \mathcal{R}_{t+\epsilon}(z) \middle| \mathcal{F}_t \right] + E \left[\int_0^\epsilon \exp(-\delta u) z_{t+u} \frac{|h_{t+u}|^2}{2} du \middle| \mathcal{F}_t \right]. \quad (42)$$

We now wish to view $r_t \doteq \mathcal{R}_t(z)$ as state variable in our decision problem. Consider first the product process zr . Then $z_t r_t$ is just the expected discounted value of a *dividend* $z_{t+u} |h_{t+u}|^2 / 2$ with discount rate δ . As such this process has a time t drift $\delta z_t r_t - z_t |h_t|^2 / 2$. In the recursive, continuous-time specification of the constraint game, the minimizing decision-maker is permitted to choose a continuation value of the conditional entropy rather than committing to entire future h process. To capture this we introduce a weighting $z_t g_t^*$ on the Brownian motion and let g_t^* be chosen by the minimizing agent. Thus we take the evolution of zr to be:

$$d(z_t r_t) = \left(\delta z_t r_t - z_t \frac{|h_t|^2}{2} \right) dt + z_t g_t^* \cdot dB_t$$

where the nonnegative martingale z evolves according to:

$$dz_t = z_t h_t \cdot dB_t.$$

Since r_t can be expressed as $(z_t r_t) / z_t$ provided that z_t is positive, an application of Ito's formula yields the first equation of (41) with $g_t = g_t^* - r_t h_t$. The second equation then follows by substituting $h_t + \tilde{B}_t$ for dB_t . Instead of having the minimizing agent choose g_t^* , we may equivalently have this agent choose g_t .

In the recursive version of the constraint game, the state variable r_t is the continuation entropy left to allocate across states at future dates. We restrict r_t to be allocated across states that can be realized with positive probability, conditional on date t information. The state variable r_t is initialized at η at date zero. The process g becomes a control vector for allocating continuation entropy across the various realized states. The vector g_t does not affect the date t local mean of the continuation entropy, but it does alter the entropy that can be allocated in the future.




Figure 1: A depiction of the recursive constraint game in discrete time using a binomial specification. The entropy level r_t is given as of date t and the probability of realizing on each of two states is one-half under the approximating model. At date t the minimizing player chooses a distorted state probability and state-dependent continuation entropy levels $r_{t+1}(\omega)$ subject to an entropy constraint. The current period r_t limits the probability distortion and next period continuation entropies. Next periods continuation entropies are discounted, averaged using the distorted probabilities and added to the relative entropy: $p_t \log p_t + (1 - p_t) \log(1 - p_t) + \log 2$.

A discrete-time analog is depicted in Figure 1. This figure considers a formulation in which two states are accessible from date t , and it is motivated by a binomial approximation to a Brownian motion. Suppose initially a move upward and downward can each occur with probability one half under the approximating model. Distorting these probabilities is achieved in our continuous-time formulation by introducing a nonzero mean to the Brownian increment. Continuation entropies are also set for each of the two states that might get realized next time period. These continuation entropies are specified at date t in our continuous-time formulation through the choice of the control vector g_t . In the discrete-time depiction, probabilities and continuation entropies in each of the two states are assigned at date t subject to an entropy constraint.

Remark 8.1. Notice that $\frac{z_{t+\tau}}{z_t}$ depicted as a function of B gives the Randon-Nikodym derivative of the distribution implied by $q_{t+\tau}$ conditioned on date t information, which we denote

$q_{t+\tau|t}$. This formula is applicable provided that z_t is not zero. We may use this conditional distribution to define:

$$\mathcal{R}_t^*(q) \doteq \int_0^\infty \exp(-\delta u) \int \log \left(\frac{dq_{t+u|t}}{dq_{t+u|t}^0} \right) dq_{t+u|t} du.$$

For each t and q , this conditional entropy measure is a nonnegative random variable defined on $(\Omega^*, \mathcal{F}^*)$. As with the unconditional counterpart, we may show that

$$\mathcal{R}_t^*(q)(B) = \mathcal{R}_t(z)$$

where the dependence on the Brownian motion B makes the left-hand side a nonnegative random variable on (Ω, \mathcal{F}) .

8.3 State Variable Degeneracy

Because the multiplier and constraint problems have identical outcomes and equilibrium representations, and because the (c_t, h_t) that solve game (30) (and therefore (29)) are functions of x_t alone, it must be true that the state r_t fails to influence either c_t or h_t in the equilibrium of game (40). This subsection verifies and explains the lack of dependence on r_t of these decisions in the equilibrium (40).

We reconsider equation (37)

$$V_r^*(\check{x}, \check{r}) = -\theta,$$

and verify that it holds for the solution path to constraint game (40) for a fixed θ . Construct

$$\lambda(\check{x}, \check{r}) = V_r^*(\check{x}, \check{r}).$$

From the g solution to game (40), it follows that

$$\lambda_r(\check{x}, \check{r})\check{g}^* + \lambda_x(\check{x}, \check{r})\sigma(\check{c}^*, \check{x}) = 0,$$

implying that $\phi(x_t, r_t)$ has a zero loading vector on the Brownian increment dB_t . Differentiating the Hamilton-Jacobi-Bellman equation with respect to r , implies that

$$\begin{aligned} & [\mu(\check{c}^*, \check{x}) + \sigma(\check{c}^*, \check{x})\check{h}^*] \cdot \lambda_x(\check{x}, \check{r}) + \left(\delta\check{r} - \frac{\check{h}^* \cdot \check{h}^*}{2} \right) \lambda_r(\check{x}, \check{r}) + \\ & \frac{1}{2} \text{trace} \left(\begin{bmatrix} \sigma(\check{c}^*, \check{x})' & \check{g}^* \end{bmatrix} \begin{bmatrix} \lambda_{xx}(\check{x}, \check{r}) & \lambda_{xr}(\check{x}, \check{r}) \\ \lambda_{rx}(\check{x}, \check{r}) & \lambda_{rr}(\check{x}, \check{r}) \end{bmatrix} \begin{bmatrix} \sigma(\check{c}^*, \check{x}) \\ \check{g}^{*'} \end{bmatrix} \right) = 0. \end{aligned}$$

Thus the local mean or dt coefficient of $\{\lambda(x_t, r_t)\}$ is also zero. As a consequence, this process remains time invariant at the solution to the constraint game.

It may happen that $\{r_t : t \geq 0\}$ hits the zero boundary in finite time along the solution to the constraint control problem. This occurs when it is optimal to eliminate exposure to the Brownian motion risk from some date forward. Once r_t is frozen at zero, we no longer expect

equation (37) to hold. To accommodate this possibility, the state evolution for $\{r_t : t \geq 0\}$ should be stopped whenever the zero boundary is hit. From this point forward V^* should be equated to the value function for the benchmark control problem. Thus we must add an exit time and a terminal value to the specification of game (40). Exit time problems such as this in which the termination occurs when a state variable hits a boundary and a terminal value function is specified are common in the continuous time literature on stochastic control. For instance, see Fleming and Soner (1993) Chapter 4.

8.4 Discussion of the ‘Irrelevance’ of Continuation Entropy

The lack of dependence of c and h on r is reminiscent of the λ -constant or Frisch demand functions used in microeconomics. The relative entropy constraint is forward-looking, as is the intertemporal wealth constraint that faces a consumer. For convenience, the Frisch demand functions use the Lagrange multiplier λ (the marginal utility of wealth) on a wealth constraint instead of wealth itself to depict consumer demands for alternative calendar dates.¹⁷

9 Comparison of Three Decision Problems

The previous two sections offered descriptions and interpretations of three closely related decision problems. Each has an associated Bellman partial differential equation and each implies the same control law for c . The Bellman equations for the three problems are:

Risk Sensitive Control Problem:

$$\begin{aligned} \delta \tilde{V}(\tilde{x}, \theta) = & \max_{\check{c} \in \check{C}} U(\check{c}, \tilde{x}) + \mu(\check{c}, \tilde{x}) \cdot \tilde{V}_x(\tilde{x}, \theta) + \frac{1}{2} \text{trace} \left[\sigma(\check{c}, \tilde{x})' \tilde{V}_{xx}(\tilde{x}, \theta) \sigma(\check{c}, \tilde{x}) \right] \\ & - \frac{1}{2\theta} \tilde{V}_x(\tilde{x}, \theta)' \sigma(\check{c}, \tilde{x}) \sigma(\check{c}, \tilde{x})' \tilde{V}_x(\tilde{x}, \theta) \end{aligned}$$

¹⁷See Frisch (1959) for a preliminary discussion and Heckman (1974) for an initial application to an intertemporal optimization problem with time separable preferences. There are extensive formal connections between the recursivity of the constraint robust control problem and the recursivity of dynamic contracts studied by Spear and Srivastava (1987), Thomas and Worrall (1988), and Kocherlakota (1996). The recursive contract literature makes the problem of designing optimal history-dependent contracts recursive by augmenting the state to include the discounted expected utility that the contract designer promises the principal. The contract is subject to a promise-keeping constraint that takes the form of (42), with continuation utility of the principal playing the role that continuation entropy does in our problem. Subject to the promise-keeping constraint, the contract designer chooses how to make continuation utility respond to the arrival of new information. As time and chance unfold, the dependence of payments on the promised value makes them history-dependent. Relative to the recursive contracts specification, our problem has some special features that allow the time- t decisions not to be history dependent, as we have shown. These special features are also satisfied by settings with sequentially complete markets in macroeconomics, for example Lucas (1982), in which time t wealth plays the role that continuation entropy does here, and in which the Lagrange multiplier on the wealth constraint is also time invariant.

(Scaled Martingale or Perturbation) Multiplier Robust Control Problem:

$$\begin{aligned} \delta\tilde{V}(\check{x}, \theta) &= \max_{\check{c} \in \check{C}} \min_h U(\check{c}, \check{x}) + \frac{\theta}{2} \check{h} \cdot \check{h} + [\mu(\check{c}, \check{x}) + \sigma(\check{c}, \check{x})\check{h}] \cdot \tilde{V}_x(\check{x}, \theta) \\ &\quad + \frac{1}{2} \text{trace} \left[\sigma(\check{c}, \check{x})' \tilde{V}_{xx}(\check{x}, \theta) \sigma(\check{c}, \check{x}) \right] \end{aligned}$$

Constraint Robust Control Problem:

$$\begin{aligned} \delta V^*(\check{x}, \check{r}) &= \max_{\check{c} \in \check{C}} \min_{h, \check{g}} U(\check{c}, \check{x}) + [\mu(\check{c}, \check{x}) + \sigma(\check{c}, \check{x})\check{h}] \cdot V_x^*(\check{x}, \check{r}) + \left(\delta\check{r} - \frac{\check{h} \cdot \check{h}}{2} \right) \cdot V_r^*(\check{x}, \check{r}) \\ &\quad + \frac{1}{2} \text{trace} \left(\begin{bmatrix} \sigma(\check{c}, \check{x})' & \check{g} \end{bmatrix} \begin{bmatrix} V_{xx}^*(\check{x}, \check{r}) & V_{xr}^*(\check{x}, \check{r}) \\ V_{rx}^*(\check{x}, \check{r}) & V_{rr}^*(\check{x}, \check{r}) \end{bmatrix} \begin{bmatrix} \sigma(\check{c}, \check{x}) \\ \check{g}' \end{bmatrix} \right) \end{aligned}$$

The first two problems share the same value function and the same control law for c . They differ only in their interpretation. The first features an enhanced response to risk and the second an adjustment for model misspecification. The second problem has an additional control h that is used to implement the robustness adjustment. This can be either a component part of a martingale indexing a probability measure or a perturbation of the state. The second and third problems are zero-sum two-player differential games. The third problem has an additional state variable and as a consequence a different value function. Moreover, the third problem uses an additional control g to set continuation entropy for future time periods. The control laws for c and h remain the same for the two differential games. The third problem has the virtue of linking up more immediately to the min-max expected utility theory of Gilboa and Schmeidler (1989). However, it is much easier to solve the second problem because there is no need explicitly to carry along the additional control and state associated with continuation entropy.

As noted above, in the recursive version of the constraint problem, continuation entropy r_t and the control g_t associated with it play similar roles to what ‘continuation wealth’ and ‘continuation utility’ play in recursive versions of competitive equilibria and optimal contract design problems, respectively. In each case, those state variables serve to make the problem recursive. Our problem has the special feature that while it is necessary to choose continuation entropy appropriately, the optimal solution isolates decisions for (c_t, h_t) from any dependence on continuation entropy.

10 Two Preference Orderings and Their Observational Equivalence

This section uses the preceding results to study the relationship between preference orderings induced by versions of the multiplier and constraint problems (3.4), (3.5). The implied preference orderings differ but are related at the common solution to both problems, where their indifference curves are tangent.

10.1 Preference Orderings

Throughout this section, we fix the filtration $\{\mathcal{F}_t, t \geq 0\}$ and consider a space of consumption processes c that are progressively measurable. In particular the time- t component c_t must be \mathcal{F}_t -measurable. This means that we can find a progressively measurable family of functions ϕ on $(\Omega^*, \mathcal{F}^*)$ so that $c_t = \phi_t(B)$. In particular, c_t will only depend on the process B up until date t . Although before we did not stipulate that the control is consumption, throughout this section and c will depict a progressively measurable consumption process.

We consider two preference orderings. To construct them, we use an endogenous state vector s_t :

$$ds_t = \mu_s(s_t, c_t)dt, \quad (43)$$

where this differential equation can be solved uniquely for s_t given, s_0 and a process $\{c_u : 0 \leq u < t\}$. We assume that the solution is a progressively measurable process $\{s_t : t \geq 0\}$. We define preferences to be time separable in (s_t, c_t) which can be represented using a utility function U . For a given process c , we may construct a process s and a process $\{U(s_t, c_t) : t \geq 0\}$. The utility process can be written as a function of B subject to information constraints. In particular, we can construct a function ν such that

$$\nu_t(c)(B) = U(s_t, c_t)$$

where $\nu_t(c)$ only depends on the process B up until time t .

In relation to our control problems, we think of s_t as an endogenous component of the state vector x_t . Individual agents recognize this endogeneity and take it into account in their preferences. We use s_t to make preferences nonseparable over time, as in models with habit persistence.

We now define two preference orderings. One preference ordering uses the valuation function:

$$W^*(c; \eta) = \inf_{\mathcal{R}^*(q) \leq \eta} \int_0^\infty \exp(-\delta t) \left(\int \nu_t(c) dq_t \right) dt.$$

Definition 10.1. (*Constraint preference ordering*) For any two progressively measurable c and c^* , $c^* \succeq_\eta c$ if

$$W^*(c^*; \eta) \geq W^*(c; \eta).$$

The other preference ordering uses the valuation function:

$$\tilde{W}(c; \theta) = \inf_q \int_0^\infty \exp(-\delta t) \left(\int \nu_t(c) dq_t \right) dt + \theta \mathcal{R}^*(q)$$

Definition 10.2. (*Multiplier preference ordering*) For any two progressively measurable c and c^* , $c^* \succeq_\theta c$ if

$$\tilde{W}(c^*; \theta) \geq \tilde{W}(c; \theta).$$

The first preference order has the multiple-priors form justified by Gilboa and Schmeidler (1989). The second is of a form that is commonly used to compute robust decision rules and is the form that is closest to recursive utility theory. We now explore their relation.

10.2 Relation Between the Preference Orders

The two preference orderings differ. Furthermore, given η , there exists no θ that makes the two preference orderings agree. However, the Lagrange Multiplier Theorem delivers a weaker result that is very useful to us. While globally the preference orderings differ, we can relate indifference curves that pass through a given point c^* in the consumption set, *e.g.* indifference curves that pass through the solution c^* to an optimal resource allocation problem.

Use the Lagrange Multiplier Theorem to write W^* as

$$W^*(c^*; \eta^*) = \max_{\theta} \inf_q \int_0^{\infty} \exp(-\delta t) \left(\int \nu_t(c) dq_t \right) dt + \theta [\mathcal{R}(Q) - \eta^*],$$

and let θ^* denote the maximizing value of θ , which we assume to be strictly positive. Suppose that $c^* \succeq_{\eta^*} c$. Then

$$\tilde{W}(c; \theta^*) - \theta^* \eta^* \leq W^*(c; \eta^*) \leq W^*(c^*; \eta^*) = \tilde{W}(c^*; \theta^*) - \theta^* \eta^*.$$

Thus $c^* \succeq_{\theta^*} c$.

The observational equivalence results from Claims 3.6 and 3.8 apply to consumption profile c^* . The indifference curves touch but do not cross at this point. We illustrate this relation in Figure 2.

While the preferences differ, this difference is concealed along a given equilibrium trajectory of consumption and prices. The tangency of the indifference curves implies that they are supported by the same prices. Observational equivalence claims made by econometricians commonly refer to equilibrium trajectories and not to off-equilibrium aspects of the preference orders. Furthermore, although the two preference orders differ, the multiplier preferences are of interest in their own right. See Wang (2001) for an axiomatic development of entropy-based preference orders that nests a finite state counterpart to this multiplier preference order.

11 Recursivity of the Preference Orderings

This section discusses the time consistency of our two preference orders. We use the fact that solutions of both of our two robust resource allocation problems satisfy Bellman equation (30), which depicts a Markov perfect equilibrium in a two-player zero-sum game. For both the multiplier and the constraint specifications, we must describe the date $t > 0$ preferences that are consistent with this solution.

At date $t > 0$, information has been realized and some consumption has taken place. Our preference orderings focus the attention of the decision maker on current and subsequent consumption in states that can be realized given current information. To study recursivity, we formulate preferences at date t to respect this vantage point. We accomplish this by using the time t conditional counterparts to the expected discounted utility process and to the relative entropy measure.




Figure 2: Indifference curves for two preference orderings that pass through a common point. There is a single time period and two states. The utility function is logarithmic in consumption in the two states. The states are equally likely under the approximating model. Probabilities are perturbed subject to a relative entropy constraint or a penalty parameter. The solid line gives the indifference curve for the preferences defined using an entropy constraint and the dashed line gives the indifference curve for the multiplier preferences.

To pose the conditional preferences, we find it convenient to use the martingale depiction on the original probability space (Ω, \mathcal{F}, P) . At date t the decision-maker cares only about states that can be realized from date t forward. That means that expectations used to average over states should be conditioned on date t information. It would be inappropriate to use date zero relative entropy to constrain probabilities conditioned on time t information.¹⁸ This leads us to use the conditional counterpart to our relative entropy measure. The date t counterpart to the multiplier preferences are based on the valuation function \tilde{W}_t given by:

$$z_t \tilde{W}_t(c; s_t \theta) = \inf_{h \in H} \int_0^\infty \exp(-\delta u) E \left(z_{t+u} U(c_{t+u}, s_{t+u}) + \theta z_{t+u} \frac{|h_{t+u}|^2}{2} \middle| \mathcal{F}_t \right) du \quad (44)$$

subject to:

$$\begin{aligned} dz_{t+u} &= z_{t+u} h_{t+u} dB_{t+u} \\ ds_{t+u} &= \mu_s(s_{t+u}, c_{t+u}) du. \end{aligned} \quad (45)$$

For the constraint preferences, at date t we ask the decision-maker to explore changes only in beliefs that affect outcomes that can be realized in the future. That is, we impose the constraint

$$\mathcal{R}_t(z) \leq r_t$$

where r_t is a state variable inherited from the minimizing agent that limits the choice of h from date t forward. Thus our constraint preferences are defined using the valuation function:

$$z_t W_t^*(c; s_t, r_t) = \inf_{h \in H} \int_0^\infty \exp(-\delta u) E \left(z_{t+u} U(c_{t+u}, s_{t+u}) \middle| \mathcal{F}_t \right) du \quad (46)$$

subject to (45) and

$$\int_0^\infty \exp(-\delta u) E \left(z_{t+u} \frac{|h_{t+u}|^2}{2} \middle| \mathcal{F}_t \right) du \leq z_t r_t$$

11.1 Multiplier Problem Revisited

We now consider the recursive nature of the optimization problem used to construct the valuation function $\tilde{W}(\cdot; \theta)$ and its relation to problem (44). For progressively measurable c , let:

$$\tilde{W}(c, \theta) = \inf_{h \in H} E \int_0^\infty \exp(-\delta u) z_u \left[U(c_u, s_u) + \theta \frac{|h_u|^2}{2} \right] du \quad (47)$$

subject to:

$$\begin{aligned} dz_u &= z_u h_u dB_t \\ ds_u &= \mu_s(s_u, c_u) dt. \end{aligned} \quad (48)$$

¹⁸Imposing a date zero relative entropy constraint at date t would introduce a temporal inconsistency by allowing the minimizing agent at date t to put no probability distortions for events at dates $u < t$ and at states that at date t are known not to have been realized. Moreover, the minimizing agent could reduce the probabilities of realized events at date t . Instead, we want the date t decision-maker to explore probability distortions that alter outcomes only from date t forward.

The problem on the right of (47) can be decomposed into two parts. First, condition on \mathcal{F}_t and a process $\{h_u : 0 \leq u < t\}$. We want to solve for $\{h_u : u \geq t\}$. The conditioning makes the problem separate across disjoint events. Therefore, we can write the optimized objective (47) as:

$$\tilde{W}(c; \theta) = \inf_{\{h_u : 0 \leq u < t\}} E \left[\int_0^t \exp(-\delta u) z_u \left(U(c_u, s_u) + \theta \frac{|h_u|^2}{2} \right) du + \exp(-\delta t) z_t \tilde{W}_t(c; s_t, \theta) \right]$$

subject to (48). Conditional rankings based on \tilde{W}_t depend on past consumptions only through the implied date t value of s_t . They do not depend on past values of h chosen by the minimizing agent.

This preference ordering is equivalent to a particular form of the stochastic differential utility studied by Duffie and Epstein (1992). This has been shown formally by Skiadis (2001) using results from Schroder and Skiadis (1999).¹⁹ This link can be illustrated as follows. Suppose that the $\tilde{C}_t \doteq \tilde{W}_t(c; s_t, \theta)$ process has a stochastic differential representation

$$d\tilde{C}_t = \omega_t dt + \kappa_t \cdot dB_t.$$

The value function process satisfies a Bellman equation

$$\begin{aligned} \delta \tilde{C}_t &= \min_{\check{h}} U(c_t, s_t) + \theta \frac{|\check{h}|^2}{2} + \omega_t + \kappa_t \cdot \check{h} \\ &= U(c_t, s_t) + \omega_t - \frac{1}{2\theta} \kappa_t \cdot \kappa_t \end{aligned}$$

[see Theorem 5 of Skiadis (2001)]. Consistent with our discussion of recursive forms of risk-sensitive control problems, the variance multiplier is $\frac{1}{\theta}$ and does not vary with the state.²⁰

The equivalence of the multiplier preference order for robustness to a risk-adjustment of the continuation value may suggest that the latter interpretation is the valid one. However, the fact that a given preference order can be motivated for alternative reasons does not inform us as to which is the appropriate motivation. The robustness motivation would lead a calibrator to think differently about the parameter θ than the risk motivation. Moreover, link between the preference orders would vanish if we limited the concerns about model misspecification to a subset of the Brownian motions.²¹ We shall return to this point in the conclusion.

11.2 Constraint Problem Revisited

We next consider the recursive nature of the optimization problem used to construct the valuation function $W^*(\cdot; \theta)$ and its relation to problem (46). For progressively measurable c ,

¹⁹Skiadis (2001) studies a finite horizon problem without an endogenous state variable. We suspect that these differences are not essential to the conclusion.

²⁰The function used to adjust for risk in the continuation value is $-\exp(-\frac{1}{\theta}W)$.

²¹In fact in Wang (2001)'s axiomatic treatment, the preferences are defined over both the approximating model and the family of perturbed models. Both can vary. By limiting the family of perturbed models we can break the link with recursive utility theory.

let:

$$W^*(c, \eta) = \inf_{h \in H} E \int_0^\infty \exp(-\delta u) z_u U(c_u, s_u) du$$

subject to (48) and $\mathcal{R}(z) \leq \eta$. Analogous to our multiplier problem, we have the recursion:

$$W^*(c; \eta) = \inf_{\{h_u: 0 \leq u < t\}, r_t \geq 0} E \left[\int_0^t \exp(-\delta u) z_u U(c_u, s_u) du + \exp(-\delta t) z_t W_t^*(c; s_t, r_t) \right]$$

subject to (48) and to:

$$E \left[\int_0^t \exp(-\delta u) z_u \frac{|h_u|^2}{2} du + \exp(-\delta t) z_t r_t \right] \leq \eta. \quad (49)$$

By standard duality arguments, W^* and W_t^* are convex in η and r_t respectively. A standard envelope argument implies that

$$\frac{dW^*(c, \eta)}{d\eta} = \frac{dW_t^*(c; s_t, r_t)}{dr}$$

which we may use to define the corresponding penalty parameter θ .²²

Implicit in the construction of the valuation function $W^*(c, \eta)$ is a partition of relative entropy over time and across states as in (49). In contrast to the multiplier problem, we are compelled to introduce a state variable r_t for the minimizing agent in addition to s_t . This state variable is introduced so that at date t the minimizer explores changes only in beliefs that affect outcomes that can be realized in the future. We tie the hands of the date t minimizer to inherit how conditional relative entropy is to be allocated across states that have already been realized at date t . The inherited r_t is determined through previous minimization.

11.3 Time Consistency

The single relative entropy constraint

$$\mathcal{R}^*(q) \leq \eta \quad (50)$$

in the time zero problem allows for tradeoffs in allocating the distortion across time periods. As a consequence, Chen and Epstein (2000) rule it out because without further constraining the decision makers (i.e., the players in the zero-sum game), the date zero constraint cannot be used in future time periods to depict preferences. Chen and Epstein (2000) want to make the decision-maker use the *full* date zero set of models at *all* dates $\tau > 0$, allowing for appropriate conditioning. Chen and Epstein's approach thus precludes using our single

²²In the absence of differentiability at some points, this relation may be extended using subgradients. In particular, this extension may be needed if either W^* or W_t^* have finite derivatives at $\eta = 0$ or $r_t = 0$ respectively.

intertemporal entropy constraint (50).²³ Nevertheless, we have shown how to implement constraint (50) recursively by limiting the model misspecifications that the decision maker can explore at time t : we restrain the choices of the time t minimizing agent in terms of an appropriately constructed continuation entropy r_t . This continuation entropy concentrates the re-evaluation of models based on date t information but requires that the restricted set of models be consistent with the entropy allocation decided earlier. This device renders the constraint formulation recursive, but not dynamically consistent in the sense advocated by Epstein and Schneider (2001). See Hansen and Sargent (2001b) for an elaboration of these issues and a further defense of our recursive formulation.

A practical reason for wanting time consistency is that it permits dynamic programming. We have shown that the recursivity of the multiplier robust control problem is sufficient to justify dynamic programming for the robust constraint control problem. Further, we can implement the constraint problem using dynamic programming with an extended state space. In effect, the constraint preferences introduce a new state variable r_t . However, as noted above, it is much easier just to solve the multiplier control problem recursively.

Another motivation for wanting time consistency is to insure that if Arrow-Debreu date and state contingent trades are made in the initial time period there will be no desire to trade such securities at future dates.²⁴ This is also true for the constraint preferences. By construction, the shadow prices for the constraint preferences match those of the multiplier preferences and the multiplier preferences are time consistent.

12 Concluding Remarks

To use the max-min expected utility theory of Gilboa and Schmeidler (1989) for applications in macroeconomics and finance, we have turned to robust control theory for parsimonious ways of specifying a decision maker's multiple models. Empirical studies in macroeconomics and finance typically assume a unique and explicitly specified dynamic statistical model. Concerns about model misspecification naturally admit that one of a set of alternative models might instead govern the data. But how should one specify those alternative models?

Robust control as described here supplies a parsimonious (one parameter) set of alternative models with rich alternative dynamics. The approach leaves those models only vaguely specified and obtains them by perturbing the decision maker's approximating model to let its shocks to feed back on state variables arbitrarily. Among other possibilities, this allows the approximating model to miss the serial correlation of exogenous variables and also to miss the dynamics of how those exogenous variables impinge on endogenous state variables. Via statistical detection error probabilities, Anderson, Hansen, and Sargent (2000) show how the multiplier parameter or the constraint parameter in the robust control problems can be

²³Chen and Epstein (2000) are led to eliminate intertemporal tradeoffs and to replace an instant-by-instant restrictions on the vector h_t .

²⁴See Johnsen and Donaldson (1985) for a discussion of time consistency and how it relates to general equilibrium theory for dynamic economies.

used to create a set of perturbed models that are difficult to distinguish statistically from the approximating model given a sample of T time-series observations.

Our two different formulations of robust control problems lead to different preference orderings but to identical decisions, and so they have tangent indifference curves at a competitive equilibrium allocation.²⁵ Other multiplier and constraint preferences could be obtained by perturbing only a subvector of the multivariate Brownian motion. Such perturbations could capture the notion that the misspecification is concentrated in only some aspects of the stochastic dynamics. Chen and Epstein (2000) use such a specification to produce preference orderings consistent with the Ellsberg paradox; it is immediate that analogous results would hold in our formulation. While there will no longer be a connection to recursive, risk-sensitive preferences, essentially the same relations will exist between robust multiplier and constraint control problems. These relations can be established by applying the Lagrange Multiplier Theorem.

A Proofs

The following is a restatement of Claim 4.4

Claim A.1. *Suppose that q_t is absolutely continuous with respect to q_t^0 for all $0 < t < \infty$. Let z be the corresponding nonnegative martingale on (Ω, \mathcal{F}, P) with $Ez_t = 1$. Then*

$$Ez_t \mathbf{1}_{\{\int_0^t |h_s|^2 ds < \infty\}} = 1,$$

where h_t is given by (14). Moreover,

$$\int \log \frac{dq_t}{dq_t^0} dq_t = \frac{1}{2} E \int_0^t z_s |h_s|^2 ds.$$

Proof. Consider first the claim that

$$Ez_t \mathbf{1}_{\{\int_0^t |h_s|^2 ds < \infty\}} = 1,$$

Formula (11) gives us the construction for the nonnegative martingale z . This martingale solves:

$$dz_t = z_t h_t dB_t$$

Construct an increasing sequence of stopping times $\{\tau_n : n \geq 1\}$ where $\tau_n \doteq \inf\{t : z_t = \frac{1}{n}\}$ and let $\tau = \lim_n \tau_n$. The limiting stopping time can be infinite. Then $z_t = 0$ for $t \geq \tau$ and

$$z_t = z_{t \wedge \tau}$$

Form:

$$z_t^n = z_{t \wedge \tau_n}$$

which is nonnegative martingale satisfying:

$$dz_t^n = z_t^n h_t^n dB_t$$

where $h_t^n = h_t$ if $0 < t < \tau_n$ and $h_t^n = 0$ if $t \geq \tau_n$. Then

$$P \left\{ \int_0^t |h_s^n|^2 (z_s^n)^2 < \infty \right\} = 1$$

²⁵We can reinterpret the solution of a stochastic growth problem as a competitive equilibrium allocation.

and hence

$$P \left\{ \int_0^t |h_s^n|^2 ds < \infty \right\} = P \left\{ \int_0^{t \wedge \tau_n} |h_s|^2 ds < \infty \right\} = 1.$$

Taking limits as n gets large,

$$P \left\{ \int_0^{t \wedge \tau} |h_s|^2 ds < \infty \right\} = 1.$$

While it is possible that $\tau < \infty$ with positive P probability, as argued by Kabanov, Lipcer, and Sirjaev (1979)

$$\int z_t \mathbf{1}_{\{\tau < \infty\}} dP = \int_{\{z_t=0, t < \infty\}} z_t dP = 0.$$

Therefore,

$$E z_t \mathbf{1}_{\{\int_0^t |h_s|^2 ds < \infty\}} = E z_t \mathbf{1}_{\{\int_0^{t \wedge \tau} |h_s|^2 ds < \infty, \tau = \infty\}} + E z_t \mathbf{1}_{\{\int_0^t |h_s|^2 ds < \infty, \tau < \infty\}} = 1.$$

Consider next the claim that

$$\int \log \frac{dq_t}{dq_t^0} dq_t = E \int_0^t z_s |h_s|^2 ds.$$

We first suppose that

$$E \int_0^t z_s |h_s|^2 ds < \infty \tag{51}$$

Use the martingale z to construct a new probability measure \tilde{P} on (Ω, \mathcal{F}) . Then from the Girsanov Theorem [see Theorem 6.2 of Lipster and Shiryaev (2000)]

$$\tilde{B}_t = B_t - \int_0^t h_s ds$$

is a Brownian motion with respect to the filtration $\{\mathcal{F}_t : t \geq 0\}$. Moreover,

$$\tilde{E} \int_0^t |h_s|^2 ds = E \int_0^t z_s |h_s|^2 ds.$$

Write

$$\log z_t = \int_0^t h_s \cdot dB_s - \frac{1}{2} \int_0^t |h_s|^2 ds = \int_0^t h_s \cdot d\tilde{B}_s + \frac{1}{2} \int_0^t |h_s|^2 ds.$$

which is well defined under the \tilde{P} probability. Moreover,

$$\tilde{E} \int_0^t h_s \cdot d\tilde{B}_s = 0$$

and hence

$$\tilde{E} \log z_t = \frac{1}{2} \tilde{E} \int_0^t |h_s|^2 ds = \frac{1}{2} E \int_0^t z_s |h_s|^2 ds,$$

which is the desired equality. In particular, we have proved that $\int \log \frac{dq_t}{dq_t^0} dq_t$ is finite.

Next we suppose that

$$\int \log \frac{dq_t}{dq_t^0} dq_t < \infty.$$

A result due to Föllmer (1985) insures that

$$\frac{1}{2} \tilde{E} \int_0^t |h_s|^2 ds \leq \int \log \frac{dq_t}{dq_t^0} dq_t.$$

Follmer's result is directly applicable because $\int \log \frac{dq_t}{dq_t^0} dq_t$ is the same as the relative entropy of \tilde{P}_t with respect to P_t where \tilde{P}_t is the restriction of \tilde{P} to events in \mathcal{F}_t and P_t is define similarly. As a consequence, (51) is satisfied and the desired inequality follows from our previous argument.

Finally, notice that $\frac{1}{2}\tilde{E} \int_0^t |h_s|^2 ds$ is infinite if, and only if $\int \log \frac{dq_t}{dq_t^0} dq_t$ is infinite. \square

B Bayesian Interpretation

In this appendix, we justify the evolution equation (33) that we used to reinterpret a robust control process as the optimal Bayesian solution to a control problem. We also construct the value function V^b for the corresponding control problem. Our justification is admittedly heuristic. In addition to being casual about the smoothness of the value function, we do not formally establish a *Verification Theorem*. While Fleming and Souganidis (1989) provide a formal justification for the existence of some such evolution equation, they do not describe how to construct this equation in practice, and how to depict conveniently the problem as a Markov control problem. Our aim is to produce a Markov depiction with an augmented state vector.

Suppose that Bellman-Isaacs condition 5.1 is satisfied. Write two partial differential equations:

$$\begin{aligned} \delta \tilde{z} \tilde{V}(\tilde{x}, \theta) &= \max_{\tilde{c} \in \tilde{C}} \tilde{z} U(\tilde{c}, \tilde{x}) + \tilde{z} \frac{\theta}{2} \alpha_h(\tilde{x}) \cdot \alpha_h(\tilde{x}) + \mu(\tilde{c}, \tilde{x}) \cdot \tilde{V}_x(\tilde{x}, \theta) \tilde{z} \\ &\quad + \tilde{z} \frac{1}{2} \text{trace} \left[\sigma(\tilde{c}, \tilde{x})' \tilde{V}_{xx}(\tilde{x}, \theta) \sigma(\tilde{c}, \tilde{x}) \right] + \tilde{z} \alpha_h(\tilde{x}) \cdot \sigma(\tilde{c}, \tilde{x})' \tilde{V}_x(\tilde{x}, \theta) \end{aligned} \quad (52)$$

$$\begin{aligned} \delta \tilde{z} V^d(\tilde{X}, \theta) &= \tilde{z} \frac{\theta}{2} \alpha_h(\tilde{X}) \cdot \alpha_h(\tilde{X}) + \tilde{z} \mu^*(\tilde{X}) \cdot V_X^d(\tilde{X}, \theta) \\ &\quad + \tilde{z} \frac{1}{2} \text{trace} \left[\sigma^*(\tilde{X})' V_{XX}^d(\tilde{X}, \theta) \sigma^*(\tilde{X}) \right] + \tilde{z} \alpha_h(\tilde{X}) \cdot \sigma^*(\tilde{X})' V_x^d(\tilde{X}, \theta) \end{aligned} \quad (53)$$

where

$$\begin{aligned} \sigma^*(\tilde{X}) &\doteq \sigma[\alpha_c(\tilde{X}), \tilde{X}] \\ \mu^*(\tilde{X}) &\doteq \mu[\alpha_c(\tilde{X}), \tilde{X}]. \end{aligned}$$

Equation (52) is a Bellman equation for an infinite-horizon discounted control problem, and equation (53) is a Lyapunov equation for evaluation of an infinite horizon, discounted objective function. (In particular, it is proportional to the evaluation of the relative entropy as in (19), where $h_t = \alpha_h(X_t)$ and X_t satisfies (33).) Form the separable value function:

$$V^b(\tilde{x}, \tilde{X}, \theta) \doteq \tilde{V}(\tilde{x}, \theta) - V^d(\tilde{X}, \theta)$$

and subtract equation (53) from (52), dividing both sides by \tilde{z} .

$$\begin{aligned} \delta V^b(\tilde{x}, \tilde{X}, \theta) &= \max_{\tilde{c} \in \tilde{C}} U(\tilde{c}, \tilde{x}) + \frac{\theta}{2} \alpha_h(\tilde{x}) \cdot \alpha_h(\tilde{x}) - \frac{\theta}{2} \alpha_h(\tilde{X}) \cdot \phi_h(\tilde{X}) \\ &\quad + \mu(\tilde{c}, \tilde{x}) \cdot V_x^b(\tilde{x}, \tilde{X}, \theta) + \mu^*(\tilde{X}) \cdot V_X^b(\tilde{x}, \tilde{X}, \theta) \\ &\quad + \frac{1}{2} \text{trace} \left[\sigma(\tilde{c}, \tilde{x})' V_{xx}^b(\tilde{x}, \tilde{X}, \theta) \sigma(\tilde{c}, \tilde{x}) \right] + \frac{1}{2} \text{trace} \left[\sigma^*(\tilde{X})' V_{XX}^b(\tilde{x}, \tilde{X}, \theta) \sigma^*(\tilde{X}) \right] \\ &\quad + \alpha_h(\tilde{x}) \cdot \sigma(\tilde{c}, \tilde{x})' V_x^b(\tilde{x}, \tilde{X}, \theta) + \alpha_h(\tilde{X}) \cdot \sigma^*(\tilde{X})' V_X^b(\tilde{x}, \tilde{X}, \theta) \end{aligned}$$

In forming this differential equation from our previous ones, we have exploited the additively separable structure of V^b in computing first and second derivatives.

Consider this differential equation along the subspace where $x = X$. Then it may be rewritten as:

$$\begin{aligned} \delta V^b(\check{x}, \check{X}, \theta) &= \max_{\check{c} \in \check{C}} U(\check{c}, \check{x}) + \mu(\check{c}, \check{x}) \cdot V_x^b(\check{x}, \check{X}, \theta) + \mu^*(\check{x}) \cdot V_X^b(\check{x}, \check{X}, \theta) \\ &\quad + \frac{1}{2} \text{trace} [\sigma(\check{c}, \check{x})' V_{xx}^b(\check{x}, \check{X}, \theta) \sigma(\check{c}, \check{x})] + \frac{1}{2} \text{trace} [\sigma^*(\check{X})' V_{XX}^b(\check{x}, \check{X}, \theta) \sigma^*(\check{X})] \end{aligned} \quad (54)$$

$$+ \alpha_h(\check{X}) \cdot \sigma(\check{c}, \check{x})' V_x^b(\check{x}, \check{X}, \theta) + \alpha_h(\check{X}) \cdot \sigma^*(\check{X})' V_X^b(\check{x}, \check{X}, \theta) \quad (55)$$

Then $\check{z}V^b(\check{x}, \check{X}, \theta)$ is the Bellman equation for a control problem with discounted objective:

$$E \int_0^\infty \exp(-\delta t) z_t U(c_t, x_t) dt \quad (56)$$

and evolution:

$$\begin{aligned} dx_t &= \mu(c_t, x_t) dt + \sigma(c_t, x_t) dB_t \\ dz_t &= z_t \alpha_h(X_t) dB_t \\ dX_t &= \mu^*(X_t) dt + \sigma^*(X_t) dB_t. \end{aligned} \quad (57)$$

The evolution equation is just a rewriting of (33). The c that attains the right side of the Bellman equation can be depicted as $c = \xi_c(\check{x}, \check{X})$, but where $\xi_c(\check{x}, \check{x}) = \alpha_c(\check{x})$. At this solution, x and X have the same evolution equation, so that when x_0 and X_0 are initialized at the same value, $x_t = X_t$ for all t . This is true even though $\{X_t\}$ is an uncontrollable or exogenous state vector while $\{x_t\}$ can be influenced by the control process $\{c_t\}$.

In summary since z_t is initialized at one, V^b is the value function for a single-agent control problem with objective (56) and dynamic evolution equation (57). Provided that x_0 and X_0 are initialized at the same value, the processes $\{x_t\}$ and $\{X_t\}$ agree and the optimal control process satisfies $c_t = \alpha_c(x_t) = \xi_c(x_t, X_t)$.

C Martingale Solution

In this appendix we describe methods for showing that the solution for z is in fact a martingale.

Write the solution for the implied state X in (33) as:

$$X_t = \Phi_t(B)$$

where Φ is a progressively measurable process on $(\Omega^*, \mathcal{F}^*)$. In other words, $\Phi_t(B)$ only depends on the Brownian path up until time t . We will have cause to use the functions $\{\Phi_t : t \geq 0\}$ in our study of the solution to the perturbation problem.

Consider the state vector process associated with the solution to the perturbation problem. This state vector satisfies:

$$d\tilde{X}_t = \mu(\tilde{X}_t) dt + \sigma^*(\tilde{X}_t) \left[\alpha_h(\tilde{X}_t) dt + d\tilde{B}_t \right] \quad (58)$$

where $\{\tilde{B}\}$ is a Brownian motion defined on an alternative probability space, say $(\Omega, \tilde{\mathcal{F}}, \tilde{P})$. Recall that the perturbation problem penalizes the discounted expected square of h . Thus the solution to the perturbation problem satisfies:

$$\tilde{E} \int_0^t |\alpha_h(\tilde{X}_u)|^2 du < \infty \quad (59)$$

for any $t > 0$.

While the solution to (58) may not be representable in terms of the past history of \tilde{B} , it should satisfy the recursion:

$$\tilde{X}_t = \Phi_t(B^*)$$

$$B_t^* = \tilde{B}_t + \int_0^t \alpha_h(\tilde{X}_u) du.$$

We refer to this as a recursion because B_t^* itself constructed from past values of \tilde{X}_t . Since \tilde{X}_t can be expressed as a function of past B_t^* we may write

$$B_t^* = \tilde{B}_t + \int_0^t \tilde{\Phi}_u(B^*) du$$

for a progressively measurable $\tilde{\Phi}$ defined on $(\Omega^*, \mathcal{F}^*)$. Moreover, for each t

$$\tilde{P} \left\{ \int_0^t |\tilde{\Phi}_u(B^*)|^2 du < \infty \right\} = 1.$$

since inequality (59) is satisfied. It follows from Theorem 7.5 of Lipster and Shiryaev (2000) that the probability distribution induced by B^* under the solution to the perturbation problem is absolutely continuous with respect to Wiener measure q^0 .

Let κ_t denote the Radon-Nikodym derivative. Then

$$Z_t = \kappa_t(B)$$

is a nonnegative martingale defined on (Ω, \mathcal{F}, P) and is the unique solution to the stochastic differential equation:

$$dZ_t = Z_t \alpha_h(X_t) dB_t$$

subject to the initial condition $Z_0 = 1$. See Theorem 7.6 of Lipster and Shiryaev (2000).

The preceding argument used the fact that the solution to the perturbation problem satisfied inequality (59). In fact, all that is needed is the weaker requirement that

$$\tilde{P} \left\{ \int_0^t |\alpha_h(\tilde{X}_u)|^2 du < \infty \right\} = 1.$$

To explore this weaker inequality, recall that

$$\alpha_h(\tilde{x}) = -\frac{1}{\theta} \sigma^*(\tilde{x})' \tilde{V}_x(\tilde{x}, \theta).$$

Provided that σ^* and $\tilde{V}_x(\cdot, \theta)$ are continuous in \tilde{x} and that \tilde{X} does not explode in finite time, this inequality follows immediately.

Another strategy for checking absolute continuity is to follow the approach of Kunita (1969), who provides characterizations of absolute continuity in Markov models through conditions on the generators of the processes. Since the models for X and \tilde{X} are Markov, we could apply these characterizations.

References

- Anderson, E., L. P. Hansen, and T. Sargent (2000, March). Robustness, detection and the price of risk. Mimeo.
- Blackwell, D. and L. Dubins (1962). Merging of opinions with increasing information. *Annals of Mathematical Statistics* 38, 882–886.

- Blackwell, D. and M. Girshick (1954). *Theory of Games and Statistical Decisions*. New York: Wiley.
- Chamberlain, G. (2000). Econometric applications of maxmin expected utility theory. *Journal of Applied Econometrics* 15, 625–644.
- Chen, Z. and L. G. Epstein (2000). Ambiguity, risk and asset returns in continuous time. unpublished manuscript.
- Duffie, D. and L. G. Epstein (1992). Stochastic differential utility. *Econometrica* 60(2), 353–394.
- Dupuis, P. and R. S. Ellis (1997). *A Weak Convergence Approach to the Theory of Large Deviations*. Wiley Series in Probability and Statistics. New York: John Wiley and Sons.
- Ekeland, I. and T. Turnbull (1983). *Infinite-Dimensional Optimization and Convexity*. Chicago Lectures in Mathematics. Chicago: The University of Chicago Press.
- Epstein, L. and M. Schneider (2001). Recursive multiple-priors. unpublished.
- Epstein, L. and S. Zin (1989). Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica* 57, 937–969.
- Fan, K. (1952). Fixed-point and minimax theorems in locally convex topological linear spaces. *Proceedings of the National Academy of Sciences* 38, 121–126.
- Fan, K. (1953). Minimax theorems. *Proceedings of the National Academy of Sciences* 39, 42–47.
- Fleming, W. H. and H. M. Soner (1993). *Controlled Markov Processes and Viscosity Solutions*. Applications of Mathematics. New York: Springer-Verlag.
- Fleming, W. H. and P. E. Souganidis (1989). On the existence of value functions of two-player, zero sum stochastic differential games. *Indiana University Mathematics Journal* 38, 293 – 314.
- Föllmer, H. (1985). An entropy approach to the time reversal of diffusion processes. In M. Metivier and E. Pardoux (Eds.), *Stochastic Differential Systems*, Volume 69 of *Lecture Notes in Control and Information Sciences*, pp. 156–163. Springer-Verlag.
- Frisch, R. (1959). A complete scheme for computing all direct and cross demand elasticities in a model with many sectors. *Econometrica* 27, 177–196.
- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18, 141–153.
- Hansen, L. P. and T. Sargent (1995, May). Discounted linear exponential quadratic gaussian control. *IEEE Transactions on Automatic Control* 40(5), 968–971.
- Hansen, L. P., T. Sargent, and T. Tallarini (1999). Robust permanent income and pricing. *Review of Economic Studies* 66, 873–907.

- Hansen, L. P. and T. J. Sargent (2001a). *Elements of robust control and filtering for macroeconomics*. To be decided.
- Hansen, L. P. and T. J. Sargent (2001b). Time inconsistency of robust control? unpublished.
- Heckman, J. J. (1974). Effects of child-care program's on women's work effort. In T. W. Schultz (Ed.), *Economics of the Family: Marriage, Children and Human Capital*, pp. 491–518. Chicago, IL: University of Chicago Press.
- Jackson, M. O., E. Kalai, and R. Smorodninsky (1999, July). Bayesian representation of stochastic processes under learning: De finetti revisited. *Econometrica* 67(4), 875–893.
- Jacobson, D. H. (1973). Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games. *IEEE Transactions on Automatic Control* 18, 124–131.
- James, M. R. (1992). Asymptotic analysis of nonlinear stochastic risk sensitive control and differential games. *Mathematics of Control, Signals, and Systems* 5, 401–417.
- Johnsen, T. H. and J. B. Donaldson (1985). The structure of intertemporal preferences under uncertainty and time consistent plans. *Econometrica* 53, 1451–1458.
- Kabanov, J. M., R. S. Lipcer, and A. N. Sirjaev (1979). Absolute continuity and singularity of locally absolutely continuous probability distributions. I. *Math. USSR Sbornik* 1935, 631 – 696.
- Kalai, E. and E. Lerner (1993). Rational learning leads to nash equilibrium. *Econometrica* 61(5), 1019–1045.
- Karatzas, I. and S. E. Shreve (1991). *Brownian Motion and Stochastic Calculus* (Second ed.). New York: Springer-Verlag.
- Kocherlakota, N. (1996). Implications of efficient risk sharing without commitment. *Review of Economic Studies* 63(4), 595–609.
- Kunita, H. (1969). Absolute continuity of markov processes and generators. *Nogaya Mathematics Journal* 36, 1–26.
- Lipster, R. S. and A. N. Shiryaev (2000). *Statistics of Random Processes* (Second ed.), Volume I. General Theory of *Applications of Mathematics*. Berlin: Springer.
- Lucas, R. E. (1982). Interest rates and prices in a two-country world. *Journal of Monetary Economics* 10(3), 335–360.
- Luenberger, D. G. (1969). *Optimization by Vector Space Methods*. New York: Wiley.
- Petersen, I. R., M. R. James, and P. Dupuis (2000). Minimax optimal control of stochastic uncertain systems with relative entropy constraints. *IEEE Transactions on Automatic Control* 45, 398–412.
- Revuz, D. and M. Yor (1994). *Continuous Martingales and Brownian Motion* (Second ed.). New York: Springer-Verlag.

- Schroder, M. and S. Skiadis (1999). *Journal of Economic Theory*.
- Skiadis, C. (2001). Robust control and recursive utility. Department of Finance, Kellogg School of Management, Northwestern University.
- Spear, S. E. and S. Srivastava (1987). On repeated moral hazard with discounting. *Review of Economic Studies* 54(4), 599–617.
- Thomas, J. and T. Worrall (1988). Self enforcing wage contracts. *Review of Economic Studies* 55, 541–554.
- Tsirel'son, B. (1975). An example of a stochastic differential equation having no strong solution. *Theory of Probability and Application* 20, 427–430.
- Wang, T. (2001). Two classes of multi-prior preferences. unpublished.
- Whittle, P. (1981). Risk-sensitive linear/quadratic/gaussian control. *Advances in Applied Probability* 13, 764–777.