

Robustness and Uncertainty Aversion

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

August 24, 2001

Abstract

This paper connects robust control theory to the max-min expected utility model of uncertainty aversion. Max-min expected utility theory depicts preferences using multiple prior distributions. Robust control theory regards a unique controlled stochastic process as an approximation by introducing a set of perturbations to it. We link the two approaches by interpreting the perturbations in robust control as the multiple priors of the max-min expected utility theory. We use a Brownian motion information structure. We construct recursive representations of preferences and of dynamic robust decision problems.

1 Introduction

This paper links the max-min expected utility theory of Gilboa and Schmeidler (1989) to the applications of stochastic robust control theory proposed 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

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

where \mathcal{Q} is a set of probability measures over c, x , and δ is a discount rate.¹ In Gilboa and Schmeidler's theory, minimization over \mathcal{Q} is a way of representing aversion to uncertainty. Gilboa and Schmeidler's theory leaves open how to specify the set \mathcal{Q} in particular applications. We are interested in robust control theory partly because it supplies a practical way to specify \mathcal{Q} in applications.

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

¹We shall call c consumption. We include the state vector x in U to accommodate time nonseparabilities.

Criteria like (1) also appear as objective functions in robust control theory. In robust control theory, minimization over \mathcal{Q} is a device for exploring the consequences of model misspecification. Robust control theory generates \mathcal{Q} by statistically perturbing a single ‘approximating model’. The theory represents \mathcal{Q} implicitly through a positive penalty or multiplier parameter θ . 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 $\{c_t\}$, but nevertheless lead to the same decisions. We describe the senses in which both problems are 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.

While the multiplier and constraint problems give rise to identical decisions, it is much easier to work with the multiplier problem. Both problems are recursive and so the solutions of both satisfy Bellman equations. However, to make the constraint problem recursive requires augmenting the state to include a continuation value for relative entropy (a measure of allowable specification errors) and also augmenting the control set of the minimizing agent to include the 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 formulating the class of misspecifications that concern the decision maker as time unfolds.

It is often valuable to study decision rules from alternative vantage points. With this in mind, we describe two alternative interpretations of a robust control law. First, under a Bellman-Isaacs condition, there is an alternative probability specification for the Brownian motion under which the robust control law is optimal in a Bayesian sense. A *pessimistic* specification of the probability distribution supports this Bayesian interpretation.

A second interpretation maintains the Brownian motion model but changes the preferences of the decision-maker to be more *risk sensitive*. Risk-sensitive control theory, as initiated by Jacobson (1973), provides a tractable way to make the decision rules or control laws more responsive to risk through the use of an exponential adjustment to the objective of the decision-maker. There are a variety of results in the control theory literature that link risk-sensitivity and 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 recursive utility theory developed by Epstein and Zin (1989) and Duffie and Epstein (1992) 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. As a point of reference, section 2

describes a standard stochastic control problem without any concern about model misspecification. To measure the size of model misspecifications, section 3 defines relative entropy and obtains representations for it that are useful for posing robust control problems. Section 4 defines two ‘time-zero’ robust control problems, called the multiplier and constraint problems, respectively, each of which is posed as a zero-sum two-person game under mutual commitment by the decision maker and a malevolent nature to stochastic decision processes at time zero. Section 5 shows how the multiplier problem can be solved recursively, while section 6 shows how the constraint problem can be solved recursively by augmenting the state and control vectors appropriately. Section 5 states circumstances, summarized by the Bellman-Isaacs condition, under which there is tight connection between the recursive Markov equilibria and the date zero commitment equilibria. Section 7 briefly compares and discusses the Bellman equations for the multiplier problem, the constraint problem, and the risk sensitive control problem. Section 8 shows how to deduce the probability specification which render robust control processes optimal in a Bayesian sense. Section 9 defines and compares two preference relations over consumption sequences that are inspired by the multiplier and constraint robust control problems. Section 10 describes how both of these preference relations can be represented recursively. Section 11 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 then introduce a concern for robustness by having the decision maker suspect that his model is an approximation to some 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 sometimes use the shorthand notation \check{g} or $\{g_t\}$ to denote the process and g_t the time t -component of that process. The actions of the decision-maker form a stochastic process $\{c_t : t \geq 0\}$ that is progressively measurable. Among other things, this requires that the time t component c_t is \mathcal{F}_t measurable.² Let U denote an instantaneous utility function, and write the discounted objective as:

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

subject to:

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

²Progressive measurability requires that we view $\check{c} \doteq \{c_t : t \geq 0\}$ as a function of (t, ω) . For any $t \geq 0$, $\check{c} : [0, t] \times \Omega$ 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.

where x_0 is a given initial condition and \check{C} is a set of admissible control processes. Under the measure P , the stochastic process for x_t is generated by (3). Equation (3) will be the ‘approximating model’ of later sections, to which all other models in a set \mathcal{Q} are perturbations.

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

Assumption 2.1.

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

subject to (3) is finite.

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

To express a concern for robustness, we entertain a family of stochastic perturbations to the Brownian motion process. We measure these perturbations or distortions in terms of *relative entropy*.

3 Relative Entropy and Convexity

For the rest of this paper, we assume that the decision maker views (3) as an approximation, and therefore also takes into account a class of alternative models that are statistically difficult to distinguish from (3). We distort the probabilities P implied by (3) to get these other models. A perturbation replaces P by another probability measure Q . This section describes the perturbations and a measure of the magnitude of the perturbation that then will be used to formulate two robust decision problems.

We want to formulate a set of stochastic processes whose members Q can be statistically difficult to distinguish from P using a finite amount of data. We formalize ‘difficult to distinguish’ in terms of log likelihood ratios of transition probabilities. That will lead us to embrace perturbations of a type that have been used by James (1992) and Anderson, Hansen, and Sargent (2000). We perturb model (3) by altering the Brownian motion specification of the shock process $\{B_t : t \geq 0\}$. While $\{B_t : t \geq 0\}$ is a Brownian motion under P , we presume that under Q this process satisfies:

$$B_t = \hat{B}_t + \int_0^t h_s ds \tag{4}$$

where $\{\hat{B}_t : t \geq 0\}$ is a d -dimensional Brownian motion and where $\{h_t : t \geq 0\}$ is a progressively measurable drift distortion. Changes in the distribution for $\{B_t : t \geq 0\}$ will be parameterized as drift distortions to a *fixed* Brownian motion $\{\hat{B}_t : t \geq 0\}$. Thus under the measure P , these distortions will be identically zero and $\{B_t : t \geq 0\}$ will coincide.

Alternative measures Q on (Ω, \mathcal{F}) correspond to alternative specifications of $\{h_t : t \geq 0\}$. We write the distorted stochastic evolution in continuous time as:

$$dx_t = \mu(c_t, x_t)dt + \sigma(c_t, x_t)(h_t dt + d\hat{B}_t) \quad (5)$$

under the Brownian motion probability specification for $\{\hat{B}_t : t \geq 0\}$.

Next we show how a notion of absolute continuity leads us to specification (4).

3.1 Absolute Continuity

We want perturbations of (3) that would be difficult to distinguish from it without large amounts of data. To make statistical discrimination among alternative models difficult requires imposing some form of absolute continuity: the perturbed model should put positive probability on the same events as the approximating model. Absolute continuity between probability distributions allows forming relative likelihoods and therefore information-based measures of distance or discrepancy between probability distributions, including the measures of Kullback and Leibler (1951) and Chernoff (1952). These measures underlie statistical theories of model discrimination.

In comparing probability distributions over infinite sequences, i.e., stochastic processes, there are alternative ways to formulate absolute continuity and to measure discrepancy. For any probability distribution Q on \mathcal{F} , let Q_t be the corresponding probability distribution restricted to \mathcal{F}_t . Following Kabanov, Lipcer, and Sirjaev (1979), we use

Definition 3.1. *The probability distribution Q is said to be **locally absolutely continuous** with respect to P if Q_t is absolutely continuous with respect to P_t for all $t \geq 0$.*

This notion of absolute continuity is weaker than requiring Q to be absolutely continuous with respect to P .³ For instance, when Q is locally absolutely continuous with respect to P , a strong Law of Large Numbers that applies to probability measure P would not necessarily also apply to Q . Time series averages that converge almost surely under P may not converge under Q , so that Q may be distinguishable from P given an infinite amount of data.⁴ Local absolute continuity focuses instead on finite histories of data and allows us to construct likelihood ratios between models at any calendar date t .

Suppose that Q is locally absolutely continuous with respect to P . Then since Q_t is absolutely continuous with respect to P_t , there exists some nonnegative random variable g_t that is \mathcal{F}_t measurable and that allows us to write

$$E_Q g_t = E_P g_t q_t \quad (6)$$

³Kabanov, Lipcer, and Sirjaev (1979) define local absolute continuity through the use of stopping times, but argue that their definition is equivalent to this “simpler one”.

⁴Our 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.

for any bounded g_t that is \mathcal{F}_t measurable. The random variable q_t is called a Radon-Nikodym derivative.

To depict the change in measure as information accrues, we require a stochastic process $\{q_t : t \geq 0\}$ adapted to $\{\mathcal{F}_t\}$. As we change the date of the information (sigma algebra) \mathcal{F}_t , we obtain different but related random variables q_t . We obtain a powerful characterization of $\{q_t\}$ by first writing a time $t + \tau$ version of (6):

$$E_Q g_{t+\tau} = E_P g_{t+\tau} q_{t+\tau}.$$

Since any random variable that is \mathcal{F}_t measurable is also $\mathcal{F}_{t+\tau}$ measurable, it follows from this equation that

$$E_Q g_t = E_P g_t q_{t+\tau}.$$

The only way that this equation and (6) can both hold for all random variables g_t is if

$$E_P(q_{t+\tau} | \mathcal{F}_t) = q_t.$$

Therefore, $\{q_t : t \geq 0\}$ must be a nonnegative martingale with $q_0 = 1$. This martingale can be zero at some date, but if it attains zero with positive probability, it must remain there. Furthermore, since $\{q_t : t \geq 0\}$ is a martingale with respect to the filtration $\{\mathcal{F}_t : t \geq 0\}$ generated by a Brownian motion, it is continuous.⁵

We now provide a main result that shows how a locally absolutely continuous distribution Q can be represented using a progressively measurable drift process $\{h_t : t \geq 0\}$ with the same dimension as the Brownian motion, provided that $\{h_t : t \geq 0\}$ is square integrable over finite time horizons under Q . The claim also asserts the converse, that if we impose the same square integrability condition on a process $\{h_t : t \geq 0\}$, we can construct a measure that preserves local absolute continuity. The claim describes how a drift distortion can be disguised in a statistical sense within a multivariate Brownian motion.

Claim 3.2. *Suppose that a measure Q is locally absolutely continuous with respect to the measure P that is represented by (3). Then outside a set of Q -measure zero, the distribution Q has a representation of the same form as (3), but where the Brownian increment dB_t is replaced by $dB_t = h_t dt + d\hat{B}_t$ for $t \geq 0$, where now $d\hat{B}_t$ is a Brownian increment. The process $\{h_t : t \geq 0\}$ is progressively measurable and satisfies $Q\{\int_0^t |h_u|^2 du < \infty\} = 1$ for all $t \geq 0$. Conversely, for any $\{h_t : t \geq 0\}$ process satisfying these same regularity conditions, the distribution Q constructed by setting $dB_t = h_t dt + d\hat{B}_t$ is locally absolutely continuous with respect to P .*

The proof of this claim in the appendix closely follows arguments in Kabanov, Lipcer, and Sirjaev (1979). An important part of the claim is that Q and not P is used to check the

⁵For example, see Revuz and Yor (1994), Chapter V, Theorem 3.4.

finiteness of $\int_0^t |h_u|^2 du < \infty$. Using Q will turn out to be natural in our setting; in fact $\int_0^t |h_u|^2 du$ can be infinite with positive probability under the P measure.⁶

The appendix shows that q_t has the representation

$$q_t = \exp \left[\int_0^t h_u \cdot d\hat{B}_u - \frac{1}{2} \int_0^t |h_u|^2 du \right], \quad (7)$$

which we shall use in the next section.

3.2 Relative Entropy of a Stochastic Process

Local absolute continuity allows us to define the *relative entropy* of a stochastic process, a key concept for formulating robust control problems. Consider a scalar stochastic process $\{g_t\}$ that is progressively measurable. This process is a random variable on the following conveniently chosen product space. Form $\Omega^* = \Omega \times \mathbb{R}^+$ where \mathbb{R}^+ is the nonnegative real line. Form the corresponding sigma algebra \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 P^* as the product measure $P \times M$ where M is exponentially distributed with density $\delta \exp(-\delta t)$. We let E^* denote the expectation operator on the product space. The expectation E^* of the stochastic process $\{g_t\}$ is by construction

$$E^*(\check{g}) = \delta \int_0^\infty \exp(-\delta t) E(g_t) dt,$$

which is an exponential average over time of $E(g_t)$.

We extend this construction by using the probability measure Q defined earlier. We form $Q^* = Q \times M$. The process $\{q_t : t \geq 0\}$ is a Radon-Nikodym derivative for Q^* with respect to P^* :

$$E_{Q^*}(\check{g}) = \delta \int_0^\infty \exp(-\delta t) E(q_t g_t) dt.$$

In particular, Q^* can be used to evaluate expected discounted utility under an absolutely continuous change in measure.

We measure the discrepancy between the distributions of P and Q as the *relative entropy* between Q^* and P^* :

$$\mathcal{R}(Q) = \delta \int_0^\infty \exp(-\delta t) E_Q(\log q_t) dt$$

⁶A commonly invoked sufficient condition for $\{q_t : t \geq 0\}$ in (7) to be a martingale is the Novikov condition:

$$E_P \left[\exp \left(\int_0^t \frac{|h_u|^2}{2} du \right) \right] < \infty$$

for all $0 \leq t < \infty$. We will use Claim 3.2 instead of the Novikov condition in what follows.

$$\begin{aligned}
&= \delta \int_0^\infty \exp(-\delta t) E_Q \left(\int_0^t h_\tau \cdot d\hat{B}_\tau - \int_0^t \frac{|h_\tau|^2}{2} d\tau \right) dt \\
&= \delta \int_0^\infty E_Q \left[\exp(-\delta t) \int_0^t \frac{|h_\tau|^2}{2} d\tau \right] dt \\
&= \delta \int_0^\infty E_Q \left(\frac{|h_u|^2}{2} \right) \int_u^\infty \exp(-\delta t) dt du \\
&= \int_0^\infty \exp(-\delta u) E_Q \left(\frac{|h_u|^2}{2} \right) du.
\end{aligned} \tag{8}$$

Relative entropy is convex in the measure Q^* , *e.g.* see Dupuis and Ellis (1997).⁷ This follows from the concavity of the function $\log y$ for positive values of y . Relative entropy is nonnegative and zero only when the probability distributions P^* and Q^* agree, which for the representation under Claim 3.2 is true only when the process $\{h_t : t \geq 0\}$ is zero.

Notice that if $\mathcal{R}(Q)$ is finite, then

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

for each positive t . It follows from Claim 3.2 that if $\mathcal{R}(Q)$ is finite, Q is locally absolutely continuous with respect to P . If in addition,

$$Q \left\{ \int_0^\infty |h_u|^2 du < \infty \right\} = 1 \tag{9}$$

then Q is absolutely continuous with respect to P on $\bigvee_{t \geq 0} \mathcal{F}_t$ (see Theorem 1 of Kabanov, Lipcer, and Sirjaev (1979)). Here $\bigvee_{t \geq 0}$ denotes the smallest sigma algebra containing all of the sigma fields $\mathcal{F}_t, t \geq 0$.

We use local absolute continuity as a device to explore alternatives to rational expectations. Though the condition $\mathcal{R}(Q) < \infty$ implies *local* absolute continuity of Q with respect to P , it does not imply absolute continuity on $\bigvee_{t \geq 0} \mathcal{F}_t$. Absolute continuity requires that whenever P assigns probability zero or one to all tail events, Q must assign the same probability to these events. If a time series average converges almost surely under P , this average would also converge under Q . In this sense, disagreement between Q and P would vanish with the passage of time. Roughly speaking, two probability measures Q and P that are mutually absolutely continuous cannot be distinguished with complete confidence from infinite amounts of data because they agree about which tail events receive probability zero and one. Two processes that are mutually *locally* absolutely continuous but not absolutely continuous can be difficult to distinguish from finite amounts of data, even though with infinite data they could be. The discounted problems that we study direct 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

⁷Other measures of relative entropy for stochastic processes occur in the literature on large deviations. For example, see Dupuis and Ellis (1997) page 299.

problems, we study misspecifications that violate (9), but have finite entropy $\mathcal{R}(\mathcal{Q})$ as defined in (8). We deliberately expand the set of misspecifications that concern the decision maker to include ones that are difficult to detect from possibly large but finite amounts of data.

In what follows, we let \check{H} denote the set of all progressively measurable processes $\check{h} \doteq \{h_t : t \geq 0\}$ such that the implied Q has finite entropy: $\mathcal{R}(Q) < \infty$.

4 Two Robust Control Problems

We now have the vocabulary to state and study the relationships between two robust control problems. We use the characterization of locally absolutely continuous probability measures justified by Claim 3.2. Both the approximating and distorted models take the form (5). The stochastic process $\{B_t : t \geq 0\}$ is a Brownian motion under P , but under Q , $dB_t = d\hat{B}_t + h_t dt$ where $\{\hat{B}_t : t \geq 0\}$ is a Brownian motion. Thus we parameterize Q by the drift distortion $\{h_t\}$ and express the distorted state evolution equation:

$$\begin{aligned} dx_t &= \mu(c_t, x_t)dt + \sigma(c_t, x_t)dB_t \\ &= \mu(c_t, x_t)dt + \sigma(c_t, x_t)(h_t dt + d\hat{B}_t). \end{aligned} \tag{10}$$

Since $\{\hat{B}_t : t \geq 0\}$ remains a Brownian motion under all of the probability models $Q \in \mathcal{Q}$, we use \hat{E} to denote the expectation operator that integrates over Brownian motion specification for $\{\hat{B}_t : t \geq 0\}$. Based on Claim 3.2, we use progressively measurable drift distortions $\{h_t : t \geq 0\}$ to index alternative probability models.

Definition 4.1. *A multiplier robust control problem is:*

$$\tilde{J}(\theta) = \sup_{\check{c} \in \check{C}} \inf_{\check{h} \in \check{H}} \hat{E} \left[\int_0^\infty \exp(-\delta t) U(c_t, x_t) dt \right] + \theta \mathcal{R}(Q) \tag{11}$$

subject to (10).

Definition 4.2. *A constraint robust control problem is:*

$$J^*(\eta) = \sup_{\check{c} \in \check{C}} \inf_{\check{h} \in \check{H}} \hat{E} \left[\int_0^\infty \exp(-\delta t) U(c_t, x_t) dt \right] \tag{12}$$

subject to (10) and $\mathcal{R}(Q) \leq \eta$.

Note that $\mathcal{R}(Q) \leq \eta$ is a single intertemporal constraint on the entire path of distortions $\{h_t : t \geq 0\}$.

These two problems are closely related. 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 taken as self-evident throughout the literature on robust control and has been explored in the context

of a linear-quadratic control problem, informally in Hansen, Sargent, and Tallarini (1999), and formally in 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 2.1, 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. Because not all values of θ are admissible, we admit 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 the admissible values of θ .

Given an $\eta > 0$, add $-\theta\eta$ to the objective in (11). For a given value of θ this has no impact on the control law.⁸ 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.

Although we have parameterized alternative probabilities Q in terms of the implied drift distortions $\{h_t : t \geq 0\}$, it is convenient to view the objective of the constraint robust control problem as being linear in Q and the entropy measure \mathcal{R} in the constraint as being 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:

$$\sup_{\check{c} \in \check{C}} \inf_{\hat{h} \in \hat{H}} \sup_{\theta \geq 0} \hat{E} \left[\int_0^\infty \exp(-\delta t) U(c_t, x_t) dt \right] + \theta [\mathcal{R}(Q) - \eta],$$

where \hat{E} again denotes integration with respect to the distribution of \hat{B} . For a given \check{c} , the objective is convex in the probability distribution Q . As is well known, 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:

$$\sup_{\check{c} \in \check{C}} \sup_{\theta \geq \underline{\theta}} \inf_{\hat{h} \in \hat{H}} \hat{E} \left[\int_0^\infty \exp(-\delta t) U(c_t, x_t) dt \right] + \theta [\mathcal{R}(Q) - \eta]. \quad (13)$$

Unfortunately, the solution for \max_θ in (13) depends on the choice of \check{c} . We will say more about that later. In solving a robust control problem, we are most interested in the \check{c} that solves the constraint robust control problem. We can find the corresponding choice of θ by changing the order of $\sup_{\check{c}}$ and \sup_θ to obtain:

$$\sup_{\theta \geq \underline{\theta}} \sup_{\check{c} \in \check{C}} \inf_{\hat{h}} \hat{E} \left[\int_0^\infty \exp(-\delta t) U(c_t, x_t) dt \right] + \theta [\mathcal{R}(Q) - \eta].$$

⁸However, it will alter which θ results in the highest objective.

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

$$\sup_{\check{c} \in \check{C}} \inf_{\check{h} \in \check{H}} \hat{E} \left[\int_0^{\infty} \exp(-\delta t) U(c_t, x_t) dt \right] + \theta^* \mathcal{R}(Q)$$

which is the multiplier robust control problem (4.1). We can drop the term $-\theta^* \eta$ from the objective without affecting the extremizing choices of \check{c}, \check{h} because we are holding θ fixed at θ^* .

Claim 4.3. *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 $J(\check{c}^*, \eta)$ satisfy:

$$J(\check{c}^*, \eta) = \inf_{\check{h} \in \check{H}} \hat{E} \left[\int_0^{\infty} \exp(-\delta t) U(c_t^*, x_t) dt \right],$$

subject to $\mathcal{R}(Q) \leq \eta$ and (10). As argued by Luenberger (1969), $J(\check{c}^*, \eta)$ is decreasing and convex in η . Given η^* , we let θ^* be the negative of the slope of the subgradient of $J(\check{c}^*, \cdot)$ at η^* . In other words, θ^* is the absolute value of the slope of a line tangent to $J(\check{c}^*, \cdot)$ 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 $J(\check{c}^*, \cdot)$ and touches $J(\check{c}^*, \cdot)$ at a point η^* .

This argument, however, fails to account for the fact that the optimized choice of \check{c} may change as we alter η . Replacing $J(\check{c}^*, \cdot)$ by J^* from Definition 4.2 accounts for the optimization with respect to \check{c} and can sometimes be used in the construction. To study this further, consider the maximized objective \tilde{J} from Definition 4.1. Then

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

We shall study the consequences of the following assumption:

Assumption 4.4. *For any $\theta > \underline{\theta}$*

$$\begin{aligned} \tilde{J}(\theta) &= \max_{\check{c} \in \check{C}} \min_{\check{h} \in \check{H}} \hat{E} \left[\int_0^{\infty} \exp(-\delta t) U(c_t, x_t) dt \right] + \theta \mathcal{R}(Q) \\ &= \min_{\check{h} \in \check{H}} \max_{\check{c} \in \check{C}} \hat{E} \left[\int_0^{\infty} \exp(-\delta t) U(c_t, x_t) dt \right] + \theta \mathcal{R}(Q) \end{aligned}$$

where the extrema are calculated subject to constraint (10).

Both equalities presume that the maximum and minimum are attained. More generally, we expect the second equality to be replaced by \leq because the minimization is first in the initial decision problem. The next section tells how to verify Assumption 4.4 and discusses some of its ramifications.

Notice that

$$\begin{aligned}
\tilde{J}(\theta) &\doteq \max_{\check{c} \in \check{C}} \min_{\check{h} \in \check{H}} \hat{E} \left[\int_0^\infty \exp(-\delta t) U(c_t, x_t) dt \right] + \theta \mathcal{R}(Q) \\
&\leq \min_{\eta \geq 0} \max_{\check{c} \in \check{C}} J(c, \eta) + \theta \eta \\
&\leq \min_{\eta \geq 0} J^*(\eta) + \theta \eta \\
&= \min_{\check{h} \in \check{H}} \max_{\check{c} \in \check{C}} \hat{E} \left[\int_0^\infty \exp(-\delta t) U(c_t, x_t) dt \right] + \theta \mathcal{R}(Q).
\end{aligned} \tag{14}$$

When Assumption 4.4 is satisfied, all of the inequalities become equalities, and

$$\tilde{J}(\theta) = \min_{\eta \geq 0} J^*(\eta) + \theta \eta. \tag{15}$$

Earlier we showed that the function $J(\check{c}^*, \eta)$ is convex and decreasing in η . Then the function J^* in Definition 4.2 is decreasing and convex in η because it is the maximum of decreasing convex functions. Equality (15) 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} :

$$J^*(\eta) = \max_{\theta \geq \underline{\theta}} \tilde{J}(\theta) - \eta \theta \tag{16}$$

For a value of $\theta^* > \underline{\theta}$, formula (15) gives a corresponding value η^* as a solution to a maximization problem. For this η^* , formula (16) guarantees that there is a solution \check{h}^* or equivalently a Q^* with relative entropy η^* .

Claim 4.5. *Suppose that Assumption 4.4 is satisfied and that for $\theta > \underline{\theta}$, \check{c}^* is the maximizing choice of c for the multiplier robust control problem 4.1. Then that \check{c}^* also solves the constraint robust control problem 4.2 for $\eta = \eta^* = \mathcal{R}(Q^*)$ where η^* solves (15).*

Claims 4.3 and 4.5 fully describe the mapping between the magnitudes of the constraint η and the multiplier θ . They do not, however, imply that given η^* the implied θ^* is unique, nor for a given $\theta^* > \underline{\theta}$ do they imply that the implied η^* is unique. While Claim 4.5 maintains Assumption 4.4, Claim 4.3 does not. Thus, without Assumption 4.4, for some values of θ a solution pair $(\check{c}^*, \check{h}^*)$ 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. For any constraint, we can find a multiplier problem with the same solution pair.

The next section describes sufficient conditions for Assumption 4.4. Our interest in these conditions extends beyond Claim 4.5 because they are informative about when solutions to the multiplier problem are recursive.

Assumption 4.4 also helps us interpret the solution to the robust control problems. Consider the control problem:

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

subject to (10). This problem takes as given the distortion Q in the distribution of the Brownian motion. The optimal choice of a progressively measurable \check{c} will take account of this distortion but will not presume to influence it. This optimized solution for \check{c} is not altered by adding $\theta \mathcal{R}(Q)$ to the objective. Thus Assumption 4.4 allows us to depict a robust control solution to the multiplier problem as one which is supported by a particular distortion in the Brownian motion. The implied least favorable Q is valid probability distribution for the exogenous stochastic process $\{B_t : t \geq 0\}$, and \check{c} is the optimal control process given that distribution. In the language of Bayesian decision theory, we may depict \check{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,

$$\max_{\check{c} \in \check{C}} \hat{E} \left[\int_0^\infty \exp(-\delta t) U(c_t, x_t) dt \right] + \theta \mathcal{R}(Q)$$

is convex in Q . From the Legendre transform,

$$J^*(\eta) = \max_{\theta \geq \underline{\theta}} \min_{\check{h} \in \check{H}} \max_{\check{c} \in \check{C}} \hat{E} \left[\int_0^\infty \exp(-\delta t) U(c_t, x_t) dt \right] + \theta [\mathcal{R}(Q) - \eta]$$

subject to (10). The parameter θ may now be interpreted as a Lagrange multiplier on the entropy constraint and is optimized to produce the worst-case distortion process \check{h} that respects this constraint. Again we can view the optimized control process \check{c} from the innermost maximization as a Bayesian solution to the control problem.

5 Recursivity of the Multiplier Formulation

This section studies the recursivity of the multiplier robust control problem, building on a result from Fleming and Souganidis (1989). 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 multiplier game is a special case of the two-player, zero-sum, stochastic differential games studied by Fleming and Souganidis (1989) and also 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).

Multiplier problem 4.1 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 choosing distorted beliefs $\{h_t\}$ takes $\{c_t\}$ as given; and the decision maker choosing $\{c_t\}$ takes $\{h_t\}$ as given. Assumption 4.4 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 4.4.

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

We now let the initial state vary and define a value function \tilde{V} as the objective for the multiplier problem. In particular, $\tilde{J}(\theta) = \tilde{V}(x, \theta)$ provided that x is set to the date zero state. With a recursive solution, this same value function is valid in subsequent time periods. Fleming and Souganidis (1989) study when there exists a recursive solution to the multiplier problem 4.1.⁹ They use what is referred to as a Bellman-Isaacs condition to justify a recursive solution. This condition serves to render equilibrium outcomes for a date zero *commitment* game identical with those for 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 Bellman-Isaacs condition is:

⁹Fleming and Souganidis (1989) impose as restrictions that μ , σ and U are bounded, uniformly continuous and Lipschitz continuous with respect to x uniformly in c . They also require that the controls c and 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 c to make $\sigma(c, x)$ have a lower rank than d , which is the dimension of the Brownian motion.

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

$$\begin{aligned}
\delta\tilde{V}(x, \theta) &= \max_{c \in C} \min_h U(c, x) + \frac{\theta}{2} h \cdot h + [\mu(c, x) + \sigma(c, x)h] \cdot \tilde{V}_x(x, \theta) \\
&\quad + \frac{1}{2} \text{trace} \left[\sigma(c, x)' \tilde{V}_{xx}(x, \theta) \sigma(c, x) \right] \\
&= \min_h \max_{c \in C} U(c, x) + \frac{\theta}{2} h \cdot h + [\mu(c, x) + \sigma(c, x)h] \cdot \tilde{V}_x(x, \theta) \\
&\quad + \frac{1}{2} \text{trace} \left[\sigma(c, x)' \tilde{V}_{xx}(x, \theta) \sigma(c, x) \right].
\end{aligned} \tag{17}$$

where \tilde{V}_x is the vector of partial derivatives of \tilde{V} with respect to x and \tilde{V}_{xx} is the matrix of second derivatives.

The partial differential equation (17) is a Bellman equation for an infinite-horizon zero-sum two-player game. The diffusion specification makes this Bellman equation a partial differential equation. With different boundary conditions, partial differential equation (17) has multiple solutions. To choose the correct solution, the solution that is the actual value function, requires that we apply a *Verification Theorem* (e.g. see Theorem 5.1 of Fleming and Soner (1993)).

Fleming and Souganidis (1989) show that the freedom to exchange orders of maximization and minimization guarantees that the date zero commitment and the Markov perfect solution to the multiplier game coincide. This exchange of orders in the recursive specification, also implies that the orders of maximization and minimization may be changed in the date zero commitment problem as required in Assumption 4.4. As we shall now see, the exchange of order of extremization in Assumption 5.1 often can 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 be based on the first-order conditions for c and h . Define

$$\begin{aligned}
\psi(c, h, x) &\doteq U(c, x) + \frac{\theta}{2} h \cdot h + [\mu(c, x) + \sigma(c, x)h] \cdot \tilde{V}_x(x, \theta) \\
&\quad + \frac{1}{2} \text{trace} \left[\sigma(c, x)' \tilde{V}_{xx}(x, \theta) \sigma(c, x) \right],
\end{aligned} \tag{18}$$

and suppose that ψ is continuously differentiable in c . First, find a Nash equilibrium by solving:

$$\begin{aligned}
\frac{\partial \psi}{\partial c}(\tilde{c}, \tilde{h}, x) &= 0 \\
\frac{\partial \psi}{\partial h}(\tilde{c}, \tilde{h}, x) &= 0.
\end{aligned}$$

In particular, the first-order conditions for h are:

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

If a solution exists, is unique and suffices for extremiation, the Bellman-Isaacs condition is satisfied. This follows from the “chain rule.” Suppose that the min player goes first and computes h as a function of x and c :

$$\tilde{h} = -\frac{1}{\theta} \sigma(\tilde{c}, x)' \tilde{V}_x(x, \theta) \tag{19}$$

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

$$\frac{\partial \psi}{\partial c} + \frac{\partial h'}{\partial c} \frac{\partial \psi}{\partial h} = 0$$

where $\frac{\partial h}{\partial c}$ can be computed from the reaction function (19). 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 c and h separate. For instance, from (19), we see that \tilde{h} does not react to c in the minimization of h conditioned on c . Even with binding constraints on c , the Bellman-Isaacs condition (Assumption 5.1) is satisfied, provided that a solution exists for c .

5.3 Convexity

A third approach uses results of Fan (1952) and Fan (1953) and 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) may be applicable. Fan (1952) also requires the set of conditional minimizers and maximizers to be convex. We know from formula (19) that the minimizers of $\psi(c, \cdot, x)$ form a singleton set, which is convex for each c and x .¹⁰ Suppose also that the set of maximizers of $\psi(\cdot, h, x)$ is non-empty and convex for each h and x . Then again the Bellman-Isaacs condition (Assumption 5.1) is satisfied. Finally Fan (1953) does not require that the set C be a subset of a linear space, but instead requires that $\psi(\cdot, h, x)$ 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 $\psi(\cdot, h, x)$ to be upper semi-continuous. The extended control space must be a compact Hausdorff space. Provided that the additional points are not attained in optimization, we can apply Fan (1953) to verify Assumption 5.1.¹¹

¹⁰Notice that provided C is compact, we can use (19) to specify a compact set that contains the entire family of minimizers for each c in C and a given x .

¹¹Consider Theorem 2 of Fan (1953) applied to $-\psi(\cdot, \cdot, x)$. This theorem does not require compactness of choice set for h , only of the choice set for c . The theorem also does not require attainment when optimization

5.4 Risk Sensitivity

The Bellman equation in Assumption 5.1 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 R be an intertemporal return function, instead of maximizing ER (where E continues to mean mathematical expectation), risk-sensitive control theory maximizes $E[\exp(\theta^{-1}R)]$, 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 describe 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 is obtained for the risk-sensitive problem is the same as the one that arises from a stochastic robust multiplier problem. As emphasized by Anderson, Hansen, and Sargent (2000), this equivalence carries over to problems with discounting, provided that 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 (19) for h into the partial differential equation (17):

$$\begin{aligned}
\delta\tilde{V}(x, \theta) &= \max_{c \in C} \min_h U(c, x) + \frac{\theta}{2} h \cdot h + [\mu(c, x) + \sigma(c, x)h] \cdot \tilde{V}_x(x, \theta) \\
&\quad + \frac{1}{2} \text{trace} \left[\sigma(c, x)' \tilde{V}_{xx}(x, \theta) \sigma(c, x) \right] \\
&= \max_{c \in C} U(c, x) + \mu(c, x) \cdot \tilde{V}_x(x, \theta) \\
&\quad + \frac{1}{2} \text{trace} \left[\sigma(c, x)' \tilde{V}_{xx}(x, \theta) \sigma(c, x) \right] \\
&\quad - \frac{1}{2\theta} \tilde{V}_x(x, \theta)' \sigma(c, x) \sigma(c, x)' \tilde{V}_x(x, \theta)
\end{aligned} \tag{20}$$

The term

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

in Bellman equation (20) 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 (20) coincides with the Bellman equation for the benchmark control problem (2), (3), but with an additional term included:

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

This term is familiar from the analysis of continuous-time, stochastic specifications of recursive utility by Duffie and Epstein (1992). Notice that $\tilde{V}_x(x, \theta)' \sigma(x, x) d\hat{B}_t$ gives the local

is over the noncompact choice set. In our application, we can verify attainment directly.

Brownian contribution to the value function process $\{\tilde{V}(x_t, \theta) : t \geq 0\}$. The additional term to the Bellman equation is the negative of the local variance of the continuation value weighted by $\frac{1}{2\theta}$. Thus under the *risk sensitive* interpretation, there is no worry about misspecified dynamics. Instead the control objective has an extra enhancement due to risk that is captured by the local variance of the continuation value.

Duffie and Epstein (1992) refer to $\frac{1}{\theta}$ as the variance multiplier. Notice that the variance multiplier is state independent. This arises when an exponential risk adjustment is made to the continuation value, analogous to the exponential risk adjustment used elsewhere in the risk-sensitive control literature. 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.

6 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.¹²

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 .

6.1 An Alternative Bellman Equation

Our strategy will be to use (15) to link the value functions for the multiplier and constraint problems, then to deduce from the Bellman equation (17) a partial differential equation that

¹²There 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 (30), 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. See the discussion in the text below.

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 (17). We shall then interpret this new game as recursively solving our original robust constraint problem 4.2.

In section 4 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:

$$V^*(x, r) = \max_{\theta \geq \underline{\theta}} \tilde{V}(x, \theta) - r\theta \quad (21)$$

where $J^*(\eta) = V^*(x, r)$ provided that x is equal to the date zero state x_0 and r is set to the initial entropy constraint η . In subsequent 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^* .

Inverting (21) yields

$$\tilde{V}(x, \theta) = \inf_{r \geq 0} V^*(x, r) + \theta r. \quad (22)$$

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 ,

$$\frac{\partial V^*}{\partial r}(x, r) = -\theta. \quad (23)$$

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

$$\tilde{V}_x = V_x^* \quad (24)$$

and

$$\begin{aligned} \tilde{V}_{xx} &= V_{xx}^* + V_{xr}^* \frac{\partial r}{\partial x} \\ &= V_{xx}^* - \frac{V_{xr}^* V_{rx}^*}{V_{rr}^*}. \end{aligned} \quad (25)$$

The Bellman-Isaacs partial differential equation (17) for \tilde{V} implies a corresponding partial differential equation for V^* that can be deduced by using formulas (24) and (25). Notice first that

$$\begin{aligned} &\frac{1}{2} \text{trace} \left[\sigma(c, x)' \tilde{V}_{xx}(x, \theta) \sigma(c, x) \right] = \\ \min_g &\frac{1}{2} \text{trace} \left(\begin{bmatrix} \sigma(c, x)' & g \end{bmatrix} \begin{bmatrix} V_{xx}^*(x, r) & V_{xr}^*(x, r) \\ V_{rx}^*(x, r) & V_{rr}^*(x, r) \end{bmatrix} \begin{bmatrix} \sigma(c, x) \\ g' \end{bmatrix} \right). \end{aligned} \quad (26)$$

This equality follows because the solution to the minimization problem is

$$g^* = -\frac{\sigma(c, x)'V_{xr}^*(x, r)}{V_{rr}^*(x, r)} \quad (27)$$

and (25) is satisfied. Thus game (17) implies that

$$\begin{aligned} \delta V^*(x, r) &= \max_{c \in C} \min_{h, g} U(c, x) + [\mu(c, x) + \sigma(c, x)h] \cdot V_x^*(x, r) + \left(\delta r - \frac{h \cdot h}{2} \right) \cdot V_r^*(x, r) \\ &\quad + \frac{1}{2} \text{trace} \left(\begin{bmatrix} \sigma(c, x)' & g \end{bmatrix} \begin{bmatrix} V_{xx}^*(x, r) & V_{xr}^*(x, r) \\ V_{rx}^*(x, r) & V_{rr}^*(x, r) \end{bmatrix} \begin{bmatrix} \sigma(c, x) \\ g' \end{bmatrix} \right), \end{aligned} \quad (28)$$

given (23).

Equation (28) supports a recursive formulation of the constraint game. In particular, this is interpretable as a Bellman-Issacs equation with a new control g and a new state r with evolution:

$$dr_t = \left(\delta r_t - \frac{h_t \cdot h_t}{2} \right) dt + g_t \cdot d\hat{B}_t. \quad (29)$$

The control g_t is chosen by the minimizing agent. This interpretation is valid provided that we can show that (28) is satisfied along the solution trajectory for the implied game. Prior to addressing this point, the next section shows that (29) describes the evolution of the continuation value of relative entropy.

6.2 Recursivity of Relative Entropy

For $\tau > 0$, our measure of relative entropy can be represented recursively as:

$$\mathcal{R}(Q) = E_Q \left[\int_0^\tau \exp(-\delta t) \frac{|h_t|^2}{2} dt + \exp(-\delta \tau) \mathcal{R}_\tau(Q) \right], \quad (30)$$

where we define a conditional discrepancy measure \mathcal{R}_τ as

$$\begin{aligned} \mathcal{R}_\tau(Q) &= \delta \int_0^\infty \exp(-\delta t) E_Q (\log q_{t+\tau} - \log q_\tau | \mathcal{F}_\tau) dt \\ &= \int_0^\infty \exp(-\delta t) E_Q \left(\frac{|h_{t+\tau}|^2}{2} | \mathcal{F}_\tau \right) dt \\ &= \exp(\delta \tau) \int_\tau^\infty \exp(-\delta t) E_Q \left(\frac{|h_t|^2}{2} | \mathcal{F}_\tau \right) dt. \end{aligned} \quad (31)$$

Representation (30) induces the recursion:

$$\mathcal{R}_\tau(Q) = \exp(-\delta \epsilon) E_Q \left[\mathcal{R}_{\tau+\epsilon}(Q) | \mathcal{F}_\tau \right] + \int_0^\epsilon \exp(-\delta t) E_Q \left(\frac{|h_{t+\tau}|^2}{2} | \mathcal{F}_\tau \right) dt \quad (32)$$

We now construct a convenient representation for the continuous-time local evolution of relative entropy. Based on the entropy recursion (32), define the following process:

$$M_\tau \doteq \mathcal{R}_\tau(Q) - \mathcal{R}(Q) - \delta \int_0^\tau \mathcal{R}_t(Q) dt + \int_0^\tau \frac{|h_t|^2}{2} dt. \quad (33)$$

Then

$$M_{\tau+\epsilon} = M_\tau + \mathcal{R}_{\tau+\epsilon}(Q) - \mathcal{R}_\tau(Q) - \delta \int_\tau^{\tau+\epsilon} \mathcal{R}_t(Q) dt + \int_\tau^{\tau+\epsilon} \frac{|h_t|^2}{2} dt.$$

Taking conditional expectations,

$$\begin{aligned} E_Q(M_{\tau+\epsilon} | \mathcal{F}_\tau) &= M_\tau + [\exp(\delta\epsilon) - 1] \mathcal{R}_\tau(Q) \\ &\quad - \int_0^\epsilon \exp[\delta(\epsilon - t)] E_Q \left(\frac{|h_{t+\tau}|^2}{2} \middle| \mathcal{F}_\tau \right) dt \\ &\quad - \delta \int_0^\epsilon \exp(\delta t) \mathcal{R}_\tau(Q) dt \\ &\quad + \delta \int_0^\epsilon \exp(\delta t) \int_0^t \exp(-\delta s) E_Q \left(\frac{|h_{s+\tau}|^2}{2} \middle| \mathcal{F}_\tau \right) ds dt \\ &\quad + \int_0^\epsilon E \left(\frac{|h_{t+\tau}|^2}{2} \middle| \mathcal{F}_\tau \right) dt. \end{aligned}$$

Notice that

$$\delta \int_0^\epsilon \exp(\delta t) \mathcal{R}_\tau(Q) dt = [\exp(\delta\epsilon) - 1] \mathcal{R}_\tau(Q)$$

and by the Fubini Theorem,

$$\begin{aligned} \delta \int_0^\epsilon \exp(\delta t) \int_0^t \exp(-\delta s) E_Q \left(\frac{|h_{s+\tau}|^2}{2} \middle| \mathcal{F}_\tau \right) ds dt &= \\ \int_0^\epsilon (\exp[\delta(\epsilon - t)] - 1) E_Q \left(\frac{|h_{t+\tau}|^2}{2} \middle| \mathcal{F}_\tau \right) dt. \end{aligned}$$

Thus

$$E_Q(M_{\tau+\epsilon} | \mathcal{F}_\tau) = M_\tau,$$

showing that $\{M_t : t \geq 0\}$ is a continuous Q -martingale. Therefore, by Theorem 3.4 in Revuz and Yor (1994), we can represent it as a stochastic integral with respect to the Brownian motion \hat{B}_t :

$$M_\tau = \int_0^\tau g_t \cdot d\hat{B}_t, \quad (34)$$

where the process $\{g_t\}$ is progressively measurable.

Combining (33) and (34), we can write the local evolution of entropy recursively as:

$$d\mathcal{R}_t(Q) = \left[\delta\mathcal{R}_t(Q) - \frac{|h_t|^2}{2} \right] dt + g_t \cdot d\hat{B}_t. \quad (35)$$

Associated with this differential equation is the boundary condition that $\mathcal{R}_t(Q)$ must remain nonnegative. Thus, (35) is a version of (29) with $r_t = \mathcal{R}_t(Q)$.

6.3 State Variable Evolution Revisited

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. From (33) and (34), we can write the local evolution of entropy as in (29). The state variable r_t is initialized at η at date zero. The process is stopped if r_t hits the zero boundary. Once zero is hit, the continuation entropy remains at zero, although in many circumstances the zero boundary will never be hit. The process g becomes a control vector that governs the allocation of the 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.

6.4 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 (17) 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 (28). This subsection verifies and explains the lack of dependence of these decisions on r_t in the equilibrium (28).

We reconsider equation (23):

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

and verify that it holds for the solution path to constraint game (28) for a fixed θ . Construct

$$\phi(x, r) = V_r^*(x, r).$$

From the g solution to game (28), it follows that

$$\phi_r(x, r)g^* + \phi_x(x, r)\sigma(c^*, x) = 0,$$

implying that $\phi(x_t, r_t)$ has a zero loading vector on the Brownian increment $d\hat{B}_t$. Differentiating the Bellman-Isaacs equation with respect to r , implies that

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

Thus the local mean or dt coefficient of $\{\phi(x_t, r_t)\}$ 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 may occur 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 (23) 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 the 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 (28).

6.5 Discussion of ‘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.¹³

7 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:

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

Multiplier Robust Control Problem:

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

¹³See Frisch (1959) for a preliminary discussion and Heckman (1974) for an initial application to an intertemporal optimization problem with time separable preferences.

Constraint Robust Control Problem:

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

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 used to implement the robustness adjustment. The second and third problems are 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 a continuation entropy value 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.

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.

8 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.¹⁴ In that time-zero game, the decision maker’s (*i.e.*, the maximizing player’s) actions do not alter the distribution Q . This was a requirement for depicting the control problem as a Bayesian solution for a particular prior distribution. The Markov perfect equilibrium characterized by (17) 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, by using a version of the ‘big X , little x ’ trick

¹⁴See Hansen and Sargent (2001a) for a related discussion cast in terms of a discrete time version of a permanent income model.

common in macroeconomics, we can use the Markov solution to the robust multiplier game to depict the equilibrium of the time-0 commitment equilibrium recursively in which the minimizing player first chooses Q . This trick allows us to express an exogenous specification of $\{B_t : t \geq 0\}$ for which the control process from the robust multiplier game is optimal. We find a recursive version of the commitment solution by revisiting our discussion at the end of section 4.

Suppose that a progressively measurable process $\{c_t : t \geq 0\}$ is chosen optimally given a distortion process $\{h_t : t \geq 0\}$ for the Brownian motion.¹⁵ 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 $\{B_t : t \geq 0\}$ cannot be influenced by $\{c_t : t \geq 0\}$. The distorted evolution facing the decision-maker can be depicted recursively as:¹⁶

$$\begin{aligned} dx_t &= \mu(c_t, x_t)dt + \sigma(c_t, x_t)[\phi_h(X_t)dt + d\hat{B}_t] \\ dX_t &= \mu^*(X_t)dt + \sigma^*(X_t)d\hat{B}_t. \end{aligned} \tag{36}$$

Here the decision maker views x as a potentially controllable part of the state and X as an uncontrollable part. In (36), $c = \phi_c(x)$ and $h = \phi_h(x)$ is the Markov solution to the differential game. The coefficients for the evolution of X satisfy

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

By construction, the control process $\{c_t\}$ is not allowed to influence the state vector process $\{X_t\}$ in (36) and the drift distortion $h_t = \phi_h(X_t)$ depends only on this uncontrollable state vector. The solution $c = \psi_c(x, X)$ to the Markov control problem will satisfy: $\phi_c(x) = \psi_c(x, x)$. Provided that X_0 is initialized at the same value as x_0 , $X_t = x_t$ at the optimized solution because the stochastic evolution for the two state vector processes will coincide.¹⁷

Although it can be used to represent the same decisions that occur in the commitment equilibrium, it is convenient that the Markov perfect representation conserves on state variables: a single vector x is used instead of the pair (x, X) and hence is computationally more tractable. Also, the Markov perfect equilibrium simplifies verification of Assumption 4.4. As we have just shown, the Markov perfect equilibrium can be used to construct a Markov control problem with an expanded state for which the \check{c} solution implied by the Markov game is optimal.

¹⁵Under 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.

¹⁶See part B of the appendix for more details about the ‘big X , little x ’ evolution equation (36) and the associated value function.

¹⁷This interpretation of (17) isolates the distortions from the influence by the decision maker. We do not necessarily want to insist on this interpretation. 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 .

9 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 (4.1), (4.2). The implied preference orderings differ but are related at the common solution to both problems, where their indifference curves are tangent.

9.1 Preference Orderings

Throughout this section, we fix the filtration $\{\mathcal{F}_t, t \geq 0\}$ and consider a space of consumption processes \check{c} that are progressively measurable. In particular the time- t component c_t must be \mathcal{F}_t -measurable. This means that for each t , c_t can be expressed in terms of a (Borel measurable) function $f_t(B^t)$ where B^t denotes the B_s process for $0 \leq s \leq t$. Although we did not stipulate previously that the control is consumption, throughout this section and the next we will use the \check{c} notation to depict a progressively measurable consumption process.

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

$$ds_t = \mu_s(s_t, c_t)dt, \quad (37)$$

where this differential equation can be solved uniquely for s_t given s_0 and process $\{c_s : 0 \leq s < t\}$. We assume that the solution is a progressively measurable process $\{s_t : t \geq 0\}$.

We define preferences to be time-additively separable in (s_t, c_t) . Given s_0 , form

$$D(\check{c}) = \int_0^\infty \exp(-\delta t)u(s_t, c_t)dt$$

for s_t that solves the differential equation (37).

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:

$$W^*(\check{c}; \eta) = \inf_{\mathcal{R}(Q) \leq \eta} E_Q D(\check{c}).$$

Definition 9.1. (*Constraint preference ordering*) For any two progressively measurable \check{c} and \check{c}^* , $\check{c}^* \succeq_\eta \check{c}$ if

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

The other preference ordering uses the valuation function:

$$\tilde{W}(\check{c}; \theta) = \inf_Q E_Q D(\check{c}) + \theta \mathcal{R}(Q)$$

Definition 9.2. (*Multiplier preference ordering*) For any two progressively measurable \check{c} and \check{c}^* , $\check{c}^* \succeq_{\theta} \check{c}$ if

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

9.2 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 \check{c}^* in the consumption set, *e.g.* indifference curves that pass through the solution \check{c}^* to an optimal resource allocation problem.

Use the Lagrange Multiplier Theorem to write W^* :

$$W^*(\check{c}^*; \eta^*) = \max_{\theta} \inf_Q E_Q D(\check{c}^*) + \theta [\mathcal{R}(Q) - \eta^*],$$

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

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

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

The observational equivalence results from Claims 4.3 and 4.5 apply to consumption profile \check{c}^* . The indifference curves touch but do not cross at this point. We illustrate this relation in Figure 1.

While the preferences differ, this difference should not be revealed 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.

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.

10 Recursivity of the Preference Orderings

This section discusses the time consistency of our two preference orders. We use the fact that either of our two robust resource allocation problems can be solved recursively by solving

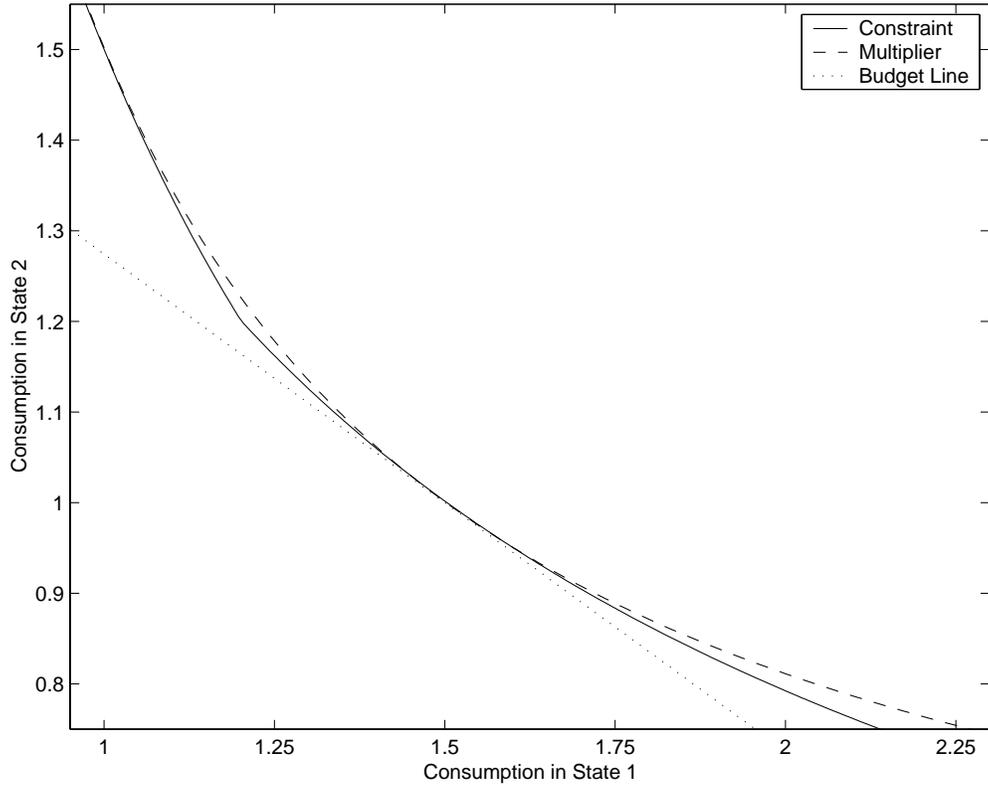


Figure 1: This figure displays indifference curves for two preference orderings that pass through a common point. The picture supposes a single time period and two states. The utility function is logarithmic in the consumption in the two states, and the states are equally likely under the approximating model. Probabilities are perturbed according to a relative entropy constraint or penalty. 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.

Bellman equation (17), 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 $\tau > 0$ preferences that are consistent with this solution.

At date $\tau > 0$, 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 τ that accommodate this change in vantage point. We accomplish this by using the time τ counterparts to D and to the relative entropy measure \mathcal{R} .

First, we construct the function:

$$D_\tau(\check{c}, s_\tau) = \int_\tau^\infty \exp(-\delta t) u(s_{t+\tau}, c_{t+\tau}) dt \quad (38)$$

where we use s_τ as the date τ initialization of the differential equation (37). The impact of consumption between date 0 and date τ is captured by the state variable s_τ . Except through s_τ , the function D_τ depends only on the consumption process from date τ forward. The function D_τ reflects a change in vantage point as time passes.

At date τ the decision maker cares only about states that can be realized from date τ forward. That means that expectations used to average over states should be conditioned on date τ information. It would be inappropriate to use date zero relative entropy to constrain probabilities conditioned on time τ information.¹⁸ This leads us to study a conditional counterpart to our relative entropy measure. Recall that our entropy measure has a recursive structure. Date zero relative entropy can be easily constructed from the conditional relative entropies in future time periods. We shall use that recursive structure in our recursive formulation of preferences. In particular, we use (38) and the recursive representation for entropy (30)-(31) to represent preferences at τ . The date τ counterpart to the multiplier preferences are based on the valuation function:

$$\tilde{W}_\tau(\check{c}; \theta) = \inf_Q E_Q [D_\tau(\check{c}, s_u) | \mathcal{F}_\tau] + \theta \mathcal{R}_\tau(Q). \quad (39)$$

For the constraint preferences, given \check{c} we can find an \tilde{h} process and an associated \tilde{Q} with which to construct $W^*(\check{c}, \eta)$. Associated with this \tilde{h} process, we can compute the time τ conditional relative entropy $\mathcal{R}_\tau(\tilde{Q})$. Thus, implicit in the construction of the valuation function $W^*(\check{c}, \eta)$ is a partition of relative entropy over time and across states as in (30). At date τ 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:

$$\mathcal{R}_\tau(Q) \leq r_\tau \quad (40)$$

¹⁸Imposing a date zero relative entropy constraint at date τ would introduce a temporal inconsistency by letting the minimizing agent put no probability distortions in dates that have already occurred and states that at date τ are known not to have been realized. Instead, we want the date τ decision-maker to explore probability distortions that alter outcomes only from date τ forward.

for $r_\tau = \mathcal{R}_\tau(\tilde{Q})$ along with fixing \tilde{h}_t for $0 \leq t < \tau$. Notice that with this constraint imposed,

$$\mathcal{R}(Q) \leq \mathcal{R}(\tilde{Q}),$$

so that we continue to satisfy our original constraint. We tie the hands of the date τ decision-maker to inherit how conditional relative entropy is to be allocated across states that have already been realized at date τ . Thus our constraint preferences are defined using the valuation function:

$$W_u(c; r_\tau) = \inf_{\mathcal{R}_\tau(Q) \leq r_\tau} E_Q [D_\tau(c, s_\tau) | \mathcal{F}_\tau]. \quad (41)$$

10.1 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 (39). For progressively measurable c , let:

$$\tilde{W}(\check{c}, \theta) = \inf_{\check{h}} \hat{E} \int_0^\infty \exp(-\delta t) \left[u(c_t, s_t) + \theta \frac{|h_t|^2}{2} \right] dt \quad (42)$$

subject to:

$$\begin{aligned} dB_t &= d\hat{B}_t + h_t dt \\ ds_t &= \mu_s(s_t, c_t) dt. \end{aligned} \quad (43)$$

The problem on the right of (42) can be decomposed into two parts. First, condition on \mathcal{F}_τ and a process $\{h_s : 0 \leq s < \tau\}$. We want to solve for $\{h_s : s \geq \tau\}$. The conditioning makes the problem separate across disjoint events. Therefore, for the first step it suffices to consider the conditional problem:

$$\inf_{\check{h}} \int_0^\infty \exp(-\delta t) \hat{E} \left[u(c_t, s_t) + \theta \frac{|h_t|^2}{2} | \mathcal{F}_u dt \right] \quad (44)$$

subject to (43) and conditioned on $\{h_t : 0 \leq t < \tau\}$. The specification of the pair $\{(c_t, h_t) : 0 \leq t < \tau\}$ has no impact on the conditional problem. In fact, we can write the optimized objective (42) as:

$$\tilde{W}(\check{c}; \theta) = \inf_{\{h_t: 0 \leq t < \tau\}} \hat{E} \left(\int_0^\tau \exp(-\delta t) \left[u(c_t, s_t) + \theta \frac{|h_t|^2}{2} \right] dt + \exp(-\delta \tau) \tilde{W}_\tau(\check{c}; \theta) \right) \quad (45)$$

subject to (43).

This preference ordering is equivalent to particular form of stochastic differential utility studied by Duffie and Epstein (1992). Let \tilde{W}_t denote the stochastic process of continuation values for a progressively measurable consumption process \check{c} . Suppose that this process has a stochastic differential representation:

$$d\tilde{W}_t = \omega_t dt + \kappa_t \cdot (h_t dt + d\hat{B}_t).$$

Although we no longer have a Markov problem (the consumption process is not necessarily a time invariant function of a Markov process), the value function process satisfies a Bellman equation:

$$\delta \tilde{W}_t = \min_h U(c_t, s_t) + \theta \frac{h^2}{2} + \omega_t + \kappa_t \cdot h.$$

Optimizing over h gives:

$$\delta \tilde{W}_t = U(c_t, s_t) + \omega_t - \frac{1}{2\theta} \kappa_t \cdot \kappa_t$$

which is a special case of equation (17) in Duffie and Epstein (1992). 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.¹⁹

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.²⁰ We shall return to this point in the conclusion.

10.2 Constraint Problem Revisited

Given $\eta = \eta^*$, let $\theta = \theta^*$ be the corresponding multiplier for the constraint problem used to construct $W^*(c, \eta)$. Substituting $\theta = \theta^*$ into the conditional problem and subtracting $\theta^* \eta^*$ gives:

$$\inf_h \int_0^\infty \exp(-\delta t) \hat{E} \left[U(c_t, s_t) + \theta^* \frac{|h_t|^2}{2} \middle| \mathcal{F}_u dt \right] - \theta^* \eta^*$$

subject to (43) and conditioned on $\{(h_t, g_t) : 0 \leq t < u\}$. We can rewrite this problem as:

$$\int_0^\tau \exp(-\delta t) U(c_t, s_t) dt + \exp(-\delta \tau) \inf_{\{h_t : t \geq \tau\}} \left(\hat{E} \left[D_\tau(c, s_\tau) + \theta^* \int_0^\infty \exp(-\delta t) \frac{|h_{t+\tau}|^2}{2} dt \middle| \mathcal{F}_\tau \right] - \theta^* r_\tau \right)$$

¹⁹The function used to adjust for risk in the continuation value is $-\exp(-\frac{1}{\theta}W)$.

²⁰In 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.

where

$$r_\tau = \exp(\delta\tau) \left[\eta^* - \int_0^\tau \exp(-\delta t) \frac{|h_t|^2}{2} dt - \int_0^\tau \exp(-\delta t) g_t \cdot d\hat{B}_t \right], \quad (46)$$

and the construction of r_τ and $\{g_t\}$ uses (33) and (34). By design, the term multiplying θ^* is zero when h is evaluated at the solution to the multiplier problem. That is, we satisfy the conditional constraint:

$$\hat{E} \left[\int_0^\infty \exp(-\delta t) \frac{|h_{t+\tau}|^2}{2} dt \middle| \mathcal{F}_\tau \right] = r_\tau$$

at the minimized choice of h . In other words, we can write the constraint problem recursively as:

$$W^*(\check{c}, \eta^*) = \inf_{\{(h_t, g_t): 0 \leq t < \tau\}} \hat{E} \int_0^\tau \exp(-\delta t) U(c_t, s_t) dt + \exp(-\delta\tau) \hat{E} W_\tau(c, r_\tau)$$

where r_τ satisfies (46). We also require r_τ to be nonnegative, which is a restriction on the admissible $\{h_t\}$ and $\{g_t\}$ processes.

10.3 Time Consistency

The single relative entropy constraint

$$\mathcal{R}(Q) \leq \eta \quad (47)$$

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 intertemporal entropy constraint (47).²¹ Nevertheless, we have shown how to implement constraint (47) recursively by limiting the model misspecifications that the decision maker can explore at time τ : we restrain the choices of the time τ minimizing agent in terms of an appropriately constructed continuation entropy r_τ . This continuation entropy concentrates the re-evaluation of models based on date τ 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 defense of our recursive formulation.

²¹Chen and Epstein (2000) are led to eliminate intertemporal tradeoffs and to replace an instant-by-instant restrictions on the vector h_t .

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 that evolves according to (29) with $r_0 = \eta$, which is consistent with the relative entropy evolution. However, 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.²² 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.

11 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 explicitly should one specify those alternative models?

Robust control supplies a parsimonious (one parameter) set of alternative models with rich alternative dynamics. The approach leaves those models 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 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.²³

Analogous multiplier and constraint preferences can 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.

²²See Johnsen and Donaldson (1985) for a discussion of time consistency and how it relates to general equilibrium theory for dynamic economies.

²³We can reinterpret the solution of a stochastic growth problem as a competitive equilibrium allocation.

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 Proof of Claim 3.2

In this appendix we provide the proof for Claim 3.2, which shows the necessity and sufficiency of representing locally absolutely continuous measures by drift distortions that satisfy certain regularity conditions. Here we break the claim into its two separate parts. First we discuss sufficiency.

Claim A.1. *Suppose that $\{h_t : t \geq 0\}$ is progressively measurable, that the distribution Q is constructed by $dB_t = h_t dt + d\hat{B}_t$ for $t \geq 0$ and that $Q\{\int_0^t |h_u|^2 du < \infty\} = 1$ for all $t \geq 0$. Then Q is locally absolutely continuous with respect to P .*

Proof. Construct an increasing sequence of stopping times

$$\tau_n = \sup\{t \leq n : \int_0^t |h_u|^2 du \leq n\}.$$

Then by the Girsanov (1960) Theorem, for each τ_n we can construct a measure Q_{τ_n} that is absolutely continuous with respect to P_{τ_n} with density:

$$q_{\tau_n} = \exp \left[\int_0^{\tau_n} h_u \cdot d\hat{B}_u - \frac{1}{2} \int_0^{\tau_n} |h_u|^2 du \right].$$

Because $\int_0^{\tau_n} |h_u|^2 du \leq n$, the Novikov condition is clearly satisfied, which guarantees the validity of this change of measure (see Karatzas and Shreve (1991)). Therefore Q is locally absolutely continuous with respect to P , when restricted to the stopping times $\{\tau_n\}$.

Then the results in Kabanov, Lipcer, and Sirjaev (1979) page 649, show that we can remove the restriction to the stopping time sequence. Let $\tau = \lim_n \tau_n$. Kabanov, Lipcer, and Sirjaev (1979) show that if κ is a stopping time, and $\lim_n Q\{\kappa > \tau_n\} = 0$, then Q_κ is absolutely continuous with respect to P_κ . In particular, since $Q\{\tau = \infty\} = 1$, we can take κ to be any $t < \infty$. Therefore Q is locally absolutely continuous with respect to P . \square

Now we prove necessity.

Claim A.2. *Suppose a measure Q is locally absolutely continuous with respect to the measure P . Then outside a set of Q -measure zero, the distribution Q has the representation of $dB_t = h_t dt + d\hat{B}_t$ for $t \geq 0$. The process $\{h_t : t \geq 0\}$ is progressively measurable and satisfies $Q\{\int_0^t |h_u|^2 du < \infty\} = 1$ for all $t \geq 0$.*

Proof. In Section 3.1, we showed that the locally absolutely continuous measure Q has a density process $\{q_t\}$ which is a continuous martingale. Next we form the continuous-time log likelihood function. Consider an increasing sequence of stopping times $\{\tau_n : n \geq 1\}$ where $\tau_n = \inf\{t : z_t < \frac{1}{n}\}$ and let $\tau = \lim_n \tau_n$. The limiting stopping time can be infinite. Following Kabanov, Lipcer, and Sirjaev (1979) construct the stochastic integral

$$M_t = \int_0^t q_u^\ominus dq_u.$$

where $q_t^\ominus = 1/q_t$ if $q_t \neq 0$ and $q_t^\ominus = 0$ if $q_t = 0$. The process $\{M_t : t \geq 0\}$ is a (τ) continuous, local martingale and can be depicted as a stochastic integral against a multivariate Brownian motion $\{\hat{B}_t : t \geq 0\}$:

$$M_t = \int_0^t h_u \cdot d\hat{B}_u$$

for $t < \tau$ where the process $\{h_t : t \geq 0\}$ is progressively measurable and satisfies:

$$\int_0^t |h_u|^2 du < \infty.$$

The restriction that $t < \tau$ guarantees that the density process has not yet reached zero. Note that it is possible that $P\{\tau < \infty\} > 0$, but we always have $Q\{\tau < \infty\} = 0$. As in Kabanov, Lipcer, and Sirjaev (1979), this follows from the local absolute continuity because:

$$Q\{\tau < \infty\} = Q\{z_t = 0, t < \infty\} = \int_{\{z_t=0, t<\infty\}} z_t dP = 0.$$

This in turn implies that $Q\{\int_0^t |h_u|^2 du < \infty\} = 1$ for all $t \geq 0$.

The Ito depiction of the log-likelihood evolution is:

$$d \log q_t = h_t \cdot d\hat{B}_t - \frac{1}{2} |h_t|^2 dt$$

for $t < \tau$. Alternatively, the exponential integral form of the density is:

$$q_t = \exp \left[\int_0^t h_u \cdot d\hat{B}_u - \frac{1}{2} \int_0^t |h_u|^2 du \right]. \quad (48)$$

Therefore, outside a set of measure zero, the Q measure has the interpretation of:

$$dB_t = h_t dt + d\hat{B}_t.$$

The Brownian motion model of $\{B_t : t \geq 0\}$ under P is replaced by a Brownian motion with a drift modelled as a progressively measurable process $\{h_t : t \geq 0\}$. This is the Girsanov (1960) Theorem (qualified with a sequence of stopping times). \square

B Bayesian Interpretation

In this appendix, we justify the evolution equation (36) 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:

$$\begin{aligned} \delta \tilde{V}(x, \theta) &= \max_{c \in C} U(c, x) + \frac{\theta}{2} \phi_h(x) \cdot \phi_h(x) + [\mu(c, x) + \sigma(c, x) \phi_h(x)] \cdot \tilde{V}_x(x, \theta) \\ &\quad + \frac{1}{2} \text{trace} \left[\sigma(c, x)' \tilde{V}_{xx}(x, \theta) \sigma(c, x) \right] \end{aligned} \quad (49)$$

$$\begin{aligned} \delta V^d(X, \theta) &= \frac{\theta}{2} \phi_h(X) \cdot \phi_h(X) + \mu^*(X) \cdot V_X^d(X, \theta) \\ &\quad + \frac{1}{2} \text{trace} \left[\sigma^*(X)' V_{XX}^d(X, \theta) \sigma^*(X) \right] \end{aligned} \quad (50)$$

where

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

Equation (49) is a Bellman equation for an infinite-horizon discounted control problem, and equation (50) 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 (8), where $h_t = \phi_h(X_t)$ and X_t satisfies (36).) Form the separable value function:

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

and subtract equation (50) from (49).

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

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:

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

This is the Bellman equation for a control problem with discounted objective:

$$\hat{E} \int_0^\infty \exp(-\delta t) U(c_t, x_t) dt \tag{52}$$

and evolution:

$$\begin{aligned} dx_t &= \mu(c_t, x_t) dt + \sigma(c_t, x_t) [\phi_h(X_t) + d\hat{B}_t] \\ dX_t &= \mu^*(X_t) dt + \sigma^*(X_t) d\hat{B}_t. \end{aligned} \tag{53}$$

The evolution equation is just a rewriting of (36). The c that attains the right side of the Bellman equation can be depicted as $c = \psi_c(x, X)$, but where $\psi_c(x, x) = \phi_c(x)$. 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, V^b is the value function for a single-agent control problem with objective (52) and dynamic evolution equation (53). 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 $c_t = \phi_c(x_t) = \psi_c(x_t, X_t)$.

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.
- Chernoff, H. (1952). A measure of the asymptotic efficiency of a hypothesis based on a sum of observations. *Annals of Mathematical Statistics* 23, 493–507.

- 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.
- 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.
- 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.
- Girsanov, I. V. (1960). On transforming a certain class of stochastic processes by absolutely continuous substitution of measures. *Theory of Probability Applications* 5, 285–301.
- 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. Smorodinsky (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.
- Kullback, S. and R. A. Leibler (1951). On information and sufficiency. *Annals of Mathematical Statistics* 22, 79–86.
- 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.
- 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.
- 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.