Time Inconsistency of Robust Control?

Lars Peter Hansen Thomas J. Sargent*

October 1, 2001

Abstract

This paper discusses senses in which alternative representations of the preferences that underlie robust control theory are or are not time consistent. The multiplier preferences of Hansen, Sargent, Turmuhambetova, and Williams (2001) are time consistent by construction. So too are their constraint preference, provided that continuation entropy is carried along as an additional state variable. Gilboa and Schmeidler's min-max expected utility theory depicts preferences using multiple prior distributions, a set of distributions that robust control theory specifies in a very parsimonious way.

1 Introduction

This paper responds to criticisms by Chen and Epstein (2000) and Epstein and Schneider (2001) of the decision theoretic foundations of robust control theory and of our work that builds on robust control theory. They focus on what they regard as an undesirable dynamic inconsistency in the preferences that robust control theorists implicitly impute to the decision maker. This paper describes various representations of robust control theory as two-player zero-sum games, provides senses of time consistency that robust control theories do and do not satisfy, and defends our opinion that the dynamic inconsistency that concerns Epstein and his coauthors is not particularly troublesome for economic applications.

Hansen, Sargent, Turmuhambetova, and Williams (2001) used ideas from robust control theory[1] to form a set of time-zero multiple priors for the min-max expected utility theory of Gilboa and Schmeidler (1989). The idea is to express the set of priors in terms of a single explicitly stated benchmark model and a constrained family of perturbations to that model. Hansen, Sargent, Turmuhambetova, and Williams (2001) call the resulting min-max preferences the *constraint preferences* because they are formulated directly in terms of a set of priors represented via a constraint on the magnitude of allowable perturbation from the

*We thank Sherwin Rosen for urging us to write this paper. We thank Nan Li and Martin Schneider for useful comments on earlier drafts.

[1]Especially Anderson, Hansen, and Sargent (2000), which builds extensively on Basar and Bernhard (1995), James (1992) and Petersen, James, and Dupuis (2000).

benchmark model. In this way, Hansen, Sargent, Turmuhambetova, and Williams (2001) connected Gilboa and Schmeidler's approach to uncertainty aversion with the literature on robust control.

As emphasized by Hansen, Sargent, Turmuhambetova, and Williams (2001), the control law that solves the time-zero robust control problem can also be expressed in terms of a recursive representation of preferences that penalizes deviations from a reference or benchmark model. These *multiplier preferences* are distinct from the date zero constraint preferences, but are related to them in convenient ways via the Lagrange Multiplier Theorem. The representation as a penalty or multiplier problem is standard in the robust control theory literature, perhaps because it is most readily conducive to computation.

The multiplier preferences used by Hansen, Sargent, and Tallarini (1999) and Anderson, Hansen, and Sargent (2000) are dynamically consistent and have been given axiomatic underpinnings by Rustichini (2000) and Wang (2001). But Chen and Epstein (2000) and Epstein and Schneider (2001) assert that the constraint preferences, which link more directly to Gilboa and Schmeidler (1989), reflect 'dynamically inconsistent' preferences. We shall argue that the type of dynamic inconsistency to which they refer differs from that familiar to macroeconomists. Indeed the constraint preferences can be depicted recursively by using an appropriate endogenous state variable. The robust control law can then be viewed as a Markov solution to a two-player, zero-sum dynamic game. As a consequence, dynamic programming methods are applicable.

This paper uses dynamic games to shed light on the concerns raised by Epstein and Schneider (2001). The representation of preferences by Gilboa and Schmeidler (1989) makes decision problems look like games. The game theoretic formulation has a long history in statistical decision theory (see Blackwell and Girshick (1954)). We will argue that the form of dynamic inconsistency that worries Epstein and his co-authors comes from halting the unfolding of the equilibrium of a two-player game midstream. Their objection amounts to a quarrel about the types of state variables that should and should not be allowed within the dynamic game used to model behavior. The state variable used in Hansen, Sargent, Turmuhambetova, and Williams (2001) arguably requires a form of commitment to the preferences orders as they are depicted in subsequent time periods.

The remainder of this paper is organized as follows. Section 2 describes Bellman equations for robust control problems. Section 3 reviews economic reasons for dynamically consistent preferences. Section 4 describes how dynamic programming applies to robust control problems. Sections 5 and 6 describe the preference orderings induced by robust control problems and alternative senses in which they are or are not time consistent. Sections 7, 8, and 9 describe the amounts of commitment, endogeneity, and separability of constraints on model misspecification built into robust control formulations, while section 10 concludes.

2 Recursive Portrayal of Robust Control Problems

A recursive version of a robust control problem can be cast in discrete time in terms of the Bellman equation

$$V(r, x) = \max_{c \in C} \min_{q^* \geq 0, r^* \geq 0} U(c, x) + \beta \int q^*(w) V[r^*(w), g(x, c, w)] F(dw)$$

where the extremization is subject to:

$$\begin{aligned} r &= \int q^*(w)[\log q^*(w) + \beta r^*(w)] F(dw) \\ 1 &= \int q^*(w) F(dw) \end{aligned}$$

In this specification, F is the distribution function for a shock vector w that is assumed to be independently and identically distributed, c is a control vector, and x is a state vector. The decision maker's approximating model asserts that next period's realized state is

$$x^* = g(x, c, w)$$

where w is the realized shock. To generate a class of perturbed models around the approximating model, the decision maker distorts the shock distribution F by using a nonnegative density q^* that serves as the Radon-Nikodym derivative of the distorted density *vis-a-vis* the benchmark approximating model.

We refer to the endogenous state variable r as conditional entropy for reasons discussed in Anderson, Hansen, and Sargent (2000). It measures the difference between two models and is related to statistical discrimination through the construction of log-likelihood ratios. The function r^* allocates next period's continuation entropy as a function of the realized shock. The pair (q^*, r^*) is constrained by the current entropy r. We assume that a discrete time Isaacs condition makes the order of minimization and maximization irrelevant.

This problem has a special structure. The envelope condition is

$$V_r(x, r) = V_r(x^*, r^*)$$

which implies a time invariant relation between x and r. As a consequence, we can depict policies that attain the right side of the Bellman equation as functions only of x: $c = \phi_c(x)$ and $q^* = \phi_q(\cdot, x)$. Moreover, it is convenient to parameterize the problem in terms of a multiplier

$$\theta = V_r(x, r)$$

that is held fixed over time. Consider instead the control problem associated with the Bellman equation:

$$W(x) = \max_{c \in C} \min_{q^* \geq 0} U(c, x) + \theta \int q^*(w) \log q^*(w) F(dw) + \beta \int q^*(w) W[g(x, c, w)] F(dw) \quad (1)$$

subject to

$$\int q^*(w)F(dw) = 1$$

This problem has one fewer state variable, implies the same solutions for q^* and c, and is more manageable computationally. Setting the multiplier θ corresponds to initializing the state variable r.

3 Why Time Consistency?

Johnsen and Donaldson (1985) contribute a valuable analysis of time consistency outside the context of model misspecification. They want a decision maker follow through with his or her initial plans as information accrues:

> Let us consider a decision maker's dynamic choice problem, as time passes and the states of the world unfold. Having carried out the current action of his chosen plan and knowing that state s obtains, he is free to choose any action in the set Y_s. Having ruled out any surprise as to what his remaining options are, if his choice deviates from the original plan, this may be taken as *prima facia* evidence of "changing tastes". If on the other hand, the original plan is carried through whatever state obtains, we may that the decision maker's tastes remain constant. His dynamic preferences will then be said to admit *time consistent planning.*

They also seek preference specifications for which there is no incentive to reopen markets at future dates provided that Arrow-Debreu contingent claims are traded at the outset. Solutions to robust control problems fulfill the *desiderata* that Johnsen and Donaldson express and produce interpretable security market price predictions.

In what follows we describe two additional time consistency issues and comment on their importance.

4 Dynamic Programming and Markov Perfect Equilibria

One reason for imposing time consistency in preferences is that it guarantees that dynamic programming methods can be applied. As we shall see, the dynamic inconsistency that concerns Chen and Epstein (2000) and Epstein and Schneider (2001) does not impede application of dynamic programming. Before discussing the kind of time inconsistency that concerns them, we briefly another time consistency issue that we view as central in robust formulations of decision problems. Robust control theorists like James (1992) and Basar and Bernhard (1995) like to emphasize the link to dynamic games. A recipe for choosing robust decisions requires a maximizing agent to rank control processes and a second malevolent agent whose distortions of probabilities relative to the approximating model induce the maximizing agent to prefer robust decisions. Thus prescriptions for robust decisions come from solving a two-player, zero-sum dynamic game (see Basar and Bernhard (1995) and James (1992)). An equilibrium of the dynamic game produces a sequence of robust decision rules. We can study how dynamic games with different timing protocols, manifested in alternative restrictions on strategies, alter equilibrium outcomes and representations.

In what follows, we use a discrete-time counterpart to the games studied by Hansen, Sargent, Turmuhambetova, and Williams (2001). Consider a two-player, zero-sum game in which one player chooses a control process $\{c_t\}$ and the other player chooses a belief process $\{q_{t+1}\}$, where q_{t+1} is nonnegative, depends on date $t+1$ information, and satisfies $E(q_{t+1}|\mathcal{F}_t) = 1$. The transition probabilities between dates t and $t+\tau$ are captured by multiplying $q_{t+1}...q_{t+\tau}$ by the τ-period transition probabilities for a benchmark model. A value process

$$V_t = U(c_t, x_t) + \beta E(q_{t+1}V_{t+1}|\mathcal{F}_t)$$

or

$$W_t = U(c_t, x_t) + E\left[q_{t+1}(\theta \log q_{t+1} + \beta W_{t+1})|\mathcal{F}_t\right]$$

can be constructed recursively, where $E(\cdot|\mathcal{F}_t)$ is the expectation operator associated with the benchmark model and $\mathcal{F}_t$ is the sigma algebra of date t events.[2] Notice that the date t recursions depend on the pair (c_t, q_{t+1}). No symptom of time inconsistency appears in these recursions. The robustness games have one player choosing c_t by maximizing and the other choosing q_{t+1} by minimizing subject to intertemporal constraints, as in the two robust decision problems described in the previous section.

Time consistency issues are resolved by verifying the Isaacs condition that guarantees that the equilibrium of the date zero commitment game coincides with the Markov perfect equilibrium. Whenever the Markov perfect equilibrium is of interest, recursive methods are known to be appropriate. The equivalence of the equilibria of these two-player zero-sum games under different timing protocols (e.g., commitment of both players to sequences at time 0 versus sequential decision making by both players) is central to the results in James (1992), Basar and Bernhard (1995), and Hansen, Sargent, Turmuhambetova, and Williams (2001).

The notion of time consistency satisfied by robust control problems is distinct from the notion of dynamic consistency that concerns Chen and Epstein (2000) and Epstein and Schneider (2001). To see the source of the difference, recall that when Gilboa and Schmeidler (1989) construct preferences that accommodate uncertainty aversion, they solve a minimization problem over measures *for each hypothetical consumption process*, instead of computing values for decision *pairs* (c_t, q_{t+1}), as in the dynamic games. A dynamic counterpart to

[2]While we have changed notation relative to that used in section 2, there is a simple relation. Since q^* was a function of w before and could be chosen to depend on x, the earlier q^* when evaluated at x_t and w_{t+1} is a $\mathcal{F}_{t+1}$ measurable random variable.

Gilboa and Schmeidler's procedure would take as a starting point a given consumption process $\{c_t\}$ and then minimize over the process $\{q_{t+1}\}$, subject to an appropriate constraint. A time consistency problem manifests itself in the solution of this problem for alternative choices of $\{c_t\}$, as we will see below. Nevertheless, the presence of this form of time consistency problem does not lead to incentives to re-open markets nor does it subvert dynamic programming.

5 A Recursive Portrayal of Preferences

Using recursions analogous to the ones described above, we can also define preferences that minimize over the process $\{q_{t+1}\}$. Suppose now that the control is consumption and that the utility function U depends only on c_t, for simplicity.[3] To define preferences we construct a value function for a general collection of consumption processes that are restricted by information constraints but are not restricted to be functions of an appropriately chosen Markov state.

We begin with a recursive constraint formulation of preferences in discrete time that uses a convenient recursive specification of a discounted version of the entropy of a stochastic process. We display it in order to understand better the sense in which the resulting preferences are recursive and to investigate their time consistency.

Given a consumption process $\{c_t : t \geq 0\}$, define

$$V_t^*(r) = \min_{q^*, r^*} U(c_t) + \beta E\left[q^* V_{t+1}^*(r^*)|\mathcal{F}_t\right] \tag{2}$$

subject to

$$\begin{aligned} r &= E\left[q^*(\log q^* + \beta r^*)|\mathcal{F}_t\right] \\ 1 &= E(q^*|\mathcal{F}_t), \end{aligned} \tag{3}$$

where now q^* and r^* are nonnegative $\mathcal{F}_{t+1}$ measurable random variables. Here we are building a function $V_t^*(\cdot)$ from $V_{t+1}^*(\cdot)$. The random variable q^* distorts the one-period transition probability. The second restriction in (3) is an adding up constraint that guarantees that the multiplication by q^* produces an alternative probability distribution.

As before, the constraint that entropy be r is used to limit the amount of model misspecification that is acknowledged, $q^* \log q^*$ is the current period contribution to entropy, and r^* is a continuation value entropy that connotes the part of entropy to be allocated in future time periods. The functions V_t^* are constructed via backward induction. The preferences are initialized using an exogenously specified value of r_0.

Holding θ fixed across alternative consumption processes gives rise to a second preference ordering. This preference-ordering can be depicted recursively, but without using entropy as

[3] Below we consider a habit persistence specification in which past consumptions are used to construct a current habit stock that enters U.

an additional state variable. The alternative recursion is

$$W_t^* = \min_{q^*} U(c_t) + \beta E\left(q^* W_{t+1}^* | \mathcal{F}_t\right) + \theta E\left(q^* \log q^* | \mathcal{F}_t\right), \tag{4}$$

which is formed as a penalty problem, where $\theta > 0$ is a penalty parameter.

Given two consumption processes, $\{c_t^1\}$ and $\{c_t^2\}$ we can construct two date zero functions $V_{0,1}^*$ and $V_{0,2}^*$ using (2) for each process. We can rank consumptions by evaluating these functions at r_0. The larger function at r_0 will tell us which of the consumption processes is preferred. For instance, if $V_{0,1}^*(r_0) \geq V_{0,2}^*(r_0)$, then the first process is preferred to the second one. Holding the penalty parameter θ fixed differs from holding fixed the entropy constraint across consumption processes, however. The value θ that makes the solution of model (4) deliver that given value of r_0 depends on the choice of the hypothesized consumption. Nevertheless, holding fixed θ gives rise to an alternative but well defined preference order. See Wang (2001) for axioms that justify these and other preferences.

6 Conditional Preference Orders

Any discussion of time inconsistency in preferences must take a stand on the preference ordering used in subsequent time periods. We now consider three different ways to construct preference orders in subsequent dates. We focus on the constraint preferences because the multiplier preferences are automatically time consistent in the sense of Johnsen and Donaldson (1985).

6.1 Implicit Preferences

Starting from date zero preferences, Johnsen and Donaldson (1985) construct an implied conditional preference order for other calendar dates, but conditioned on realized events. They then explore properties of the conditional preference order. As they emphasize, the resulting family of conditional preference orders is, by construction, time consistent. The question is whether these preference orders are appealing. To judge this, Johnsen and Donaldson (1985) define properties such as history dependence, conditional weak dependence, and dependence on unrealized alternatives.

At date zero, we can use a common r_0 to initialize the constraint preference orders. Two alternative consumption processes are, however, associated with two alternative specifications $\{q_{t+1} : 0 \leq t \leq \tau - 1\}$ and also two processes for continuation entropy r_τ. In spite of the separability over time and across states in the objective, the different choices of q_{t+1} will cause history dependence. Moreover, $V_\tau^*(r_\tau)$ in states that are known not to be realized will have an impact on the conditional preference order over states that can be realized based on date τ information. The minimization used in defining preferences apparently induces some unappealing aspects in the implied consumption rankings as time unfolds. In spite of the recursive construction, all branches in the construction of V_0^* remain relevant in reassessing the consumption preferences from the vantage point of date τ.

These aspects of the implied preference orders in subsequent time periods might seem to threaten to undermine the applicability of dynamic programming, but in fact they do not. Moreover, as we will see in section 6.3 there is another and more tractable way to specify preferences over time.

6.2 Unconstrained Reassessment of Date Zero Models

In an analysis of a continuous-time multiple priors model, Chen and Epstein (2000) take a different point of view about the intertemporal preference orders. Suppose that the date τ minimizing decision maker uses the date zero family of models but cares only about consumption from date τ forward conditioned on date τ information. The absence of dependence on past consumptions is because, at least for the moment, U depends only on c_t. Exploring the conditional probabilities implied by the full set of date zero models generates time inconsistency. The reason is as follows.

The function $V_\tau^*(\cdot)$ is constructed via backward induction. But at date τ the minimization suggested by Chen and Epstein (2000) includes minimizing over r_τ. To make the date τ conditional entropy r_τ large, the minimizing agent would make small the *ex ante* probability of the date τ observed information. For instance, suppose that τ is one. Then at date one we consider the problem:

$$\min_{q^*, r^*} V_1^*(r^*)$$

subject to:

$$\begin{aligned} r_0 &= E\left[q^*(\log q^* + \beta r^*)|\mathcal{F}_0\right] \\ 1 &= E(q^*|\mathcal{F}_0) \end{aligned}$$

where q^* and r^* are restricted to be nonnegative and $\mathcal{F}_1$ measurable. The objective is to be minimized conditioned on date one conditioning information. Notice that when q^* is zero for the realized date one information, r^* can be made arbitrarily large. Thus the date one re-optimization becomes degenerate and inconsistent with the recursive construction of V_0^*. The source of the time inconsistency is the ability of the date τ minimization to reassign distorted probabilities that apply to events that have already been realized. To avoid this problem, Chen and Epstein (2000) argue for separability in the specification of alternative models across dates and realized events. For instance, instead of the recursive constraint (3) we could require

$$\begin{aligned} E\left[q^*(\log q^*)|\mathcal{F}_t\right] &\leq \eta_t \\ E(q^*|\mathcal{F}_t) &= 1 \end{aligned} \tag{5}$$

for an exogenously specified process $\{\eta_t\}$.[4]

[4] Alternatively, Epstein and Schneider (2001) suggest that one might begin with a family of models

6.3 A Better Approach

Our recursive construction of V and V_τ^* suggests a different approach than either the implicit approach of section 6.1 or the unconstrained reassessment approach of section 6.2. Suppose that the re-optimization from date τ forward does not allow a reassessment of the distortion of probabilities of events that have already been realized as of date τ. That can be accomplished by endowing the time τ minimizing agent with a state variable r_τ. This state variable is held fixed at date τ when evaluating alternative consumption processes. We use appropriately constructed valuations $V_\tau^*(r_\tau)$ to rank consumption processes from date τ forward. Across consumption processes, the common value of the state variable r_τ was chosen earlier (as a function of date τ shocks) and is inherited by the date τ decision-maker(s). Conditioning on this state variable makes contributions from previous dates and from unrealized states irrelevant.

Distortions of the probabilities of future events that can be realized given current information can still be explored by the date τ decision-maker. Reallocation of future conditional relative entropy r^* given (3) is permitted at date τ. Given our recursive construction, this more limited type of re-exploration will not cause the preferences to be time inconsistent.

We see very little appeal to the idea of distorting probabilities of events that have already been realized, and thus are not bothered by limiting the scope of the re-evaluation in this way. Nevertheless, our formulation does require a form of commitment and a state variable to keep track of this commitment.

While this approach results in a different family of preference orders than the implicit approach, the differences are inconsequential in recursive control problems. The preferences remain consistent in the following sense. Consider the re-evaluation of the process $\{c_t^1\}$. Associated with this process is a continuation entropy r_τ for date τ. Consider an alternative process $\{c_t^2\}$ that agrees with the original process up until (but excluding) time τ. If $\{c_t^1\}$ is preferred to $\{c_t^2\}$ at date τ with probability one, then this preference ordering will be preserved at date zero.[5] The date zero problem allows for a more flexible minimization, but this flexibility will only reduce the date zero value of $\{c_t^1\}$.[6]

constrained in accordance with difference equation (3) solved forward from date zero. One could then expand this family of models sufficiently to satisfy their dynamic consistency requirement. In particular, one might hope to find an implied choice of η_t's in (5) to support this construction. Unfortunately, the attempted construction of the η_t's would suffer from an analogous problem. The restriction on the density q's in future periods would be effectively removed and the η_t's in (5) would have to be infinite. This proposed repair is uninteresting for the decision-problems described here because the expanded set of probability models is too large.

[5]This can be seen by computing a date zero value for the $\{c_t^2\}$ using the minimizing distortions between date one and τ.

[6]See also Epstein and Schneider (2001) for a closely related discussion of a weaker dynamic consistency axiom.

7 Commitment

Provided that the date τ decision-maker commits to using r_τ in ranking consumptions from date τ forward, the implied preferences by (2) are made recursive by supposing that the date τ minimizing agent can assign the continuation entropy for date $\tau+1$ chosen as a function of tomorrow's realized state. A possible complaint about this formulation is that it requires too much commitment. In ranking consumption processes from date τ forward, why should the r_τ chosen for a particular consumption process be credibly adhered to?

Some such form of commitment in individual decision-making does not seem implausible to us. We can debate how much commitment is reasonable, but then it also seems appropriate to ask what leads decision makers to commit to an exogenously specified process $\{\eta_t\}$ of entropy distortions specified period-by-period as in (5). Neither our decision-making environment nor that envisioned by Chen and Epstein (2000) and Epstein and Schneider (2001) is, in our view, rich enough to address this question.

8 Endogenous State Variable

Our representation requires an additional endogenous state variable to describe preferences. The fact that we have carried along that state variable as an argument in the function V_t^* distinguishes our formulation from usual specifications of preferences in single agent decision problems. State variables do play a role in other preference orders. For instance, preferences with intertemporal complementarities such as those with habit persistence include a state variable constructed from past consumptions called a habit stock.

To illustrate the differences between the use of a state variable to depict habit persistence and the state variable used in our preferences, suppose that the habit stock is constructed as a geometric weighted average:

$$h_t = (1-\lambda)c_t + \lambda h_{t-1}, \tag{6}$$

for $0 < \lambda < 1$. Define the date t preferences using

$$\tilde{V}_t = U(c_t, h_{t-1}) + \beta E\left(\tilde{V}_{t+1}|\mathcal{F}_t\right) \tag{7}$$

where (6) is used to build the habit stock from current and past consumption. A feature of (7) is that we may be able depict date t preferences in terms of consumption from date t forward and the habit stock h_{t-1} coming into time t. A state variable h_{t-1} is used to define the date t preferences, but this variable can be constructed mechanically from past consumption.

Consider now two consumption processes $\{c_t^1\}$ and $\{c_t^2\}$ that agree from date zero through date $\tau-1$ and suppose that h_{-1} is fixed at some arbitrary number. Thus $h_t^1 = h_t^2$ for $t = 0, 1, ..., \tau-1$. If $\tilde{V}_\tau^1 \geq \tilde{V}_\tau^2$ with probability one, then $\tilde{V}_0^1 \geq \tilde{V}_0^2$ with probability one. This is the notion of time consistency in preferences used by Duffie and Epstein (1992) and others, appropriately extended to include a state variable. Habit persistent preferences are dynamically consistent in this sense, once we introduce an appropriate a state variable into the analysis. In contrast to the conditional entropy r_t, the habit stock state variable h_{t-1} can be formed mechanically from past consumptions. There is no reference to optimization needed to construct h_{t-1} when we compare consumption processes with particular attributes.

In contradistinction, our state variable r_τ cannot be formed mechanically in terms of past consumption. It is constructed through optimization and is therefore forward-looking. This might seem unattractive because it makes the date τ preferences look 'too endogenous'. The forward-looking nature of this variable makes it depend on unrealized alternatives. (See Epstein and Schneider (2001) for an elaboration on this complaint.) Thus we are using state variables of a rather different nature than occur with time nonseparable preferences. If we condition on an initial r_0 and compare consumption processes that agree between dates zero and $\tau - 1$, we will not necessarily be led to use the same value of r_τ.

A complaint about being *too endogenous* is perhaps not so devastating. Proponents of habit persistence like to emphasize the endogeneity of the resulting preference ordering. While the habit-stock state variable can be formed mechanically, along a chosen consumption path the realized habit stock will typically depend on beliefs about the future and be forward-looking. This feature is emphasized in models of "rational addiction" and is an attribute for which apologies are not offered.[7] Whenever we have history dependence in preference orders, along a chosen consumption path the date τ preference order will depend on 'unrealized alternatives' through the endogeneity of the state variable. Just as minimization induces this dependence in our investigation, utility maximization will induce it along a chosen path. The time consistency problem in preferences over consumption processes comes from studying half of a two-player, dynamic game.

9 Separable Entropy

Our aim in studying preferences for robustness is to explore extensions of rational expectations that accommodate model misspecification. We seek convenient ways to explore the consequences of decisions across dynamic models with similar observable implications. Statistical discrimination leads us to study relative likelihoods. By their very nature, likelihood ratios involve intertemporal tradeoffs.

To accommodate misspecification in a dynamic evolution equation using a separable specification would seem to require some form state dependence in the constraints. For instance, many interesting misspecifications of a first-order autoregression would require a state-dependent restriction on the one-period conditional entropy. This state dependence is permitted by Chen and Epstein (2000) and Epstein and Schneider (2001) but its precise nature is in practice left to the researcher or decision-maker.[8] It is intractable to ex-

[7] A form of commitment is also present in habit persistent models since the date τ decision-maker remains 'committed' to past experience as measured by the habit stock $h_{\tau-1}$.

[8] Epstein and Schneider (2001) feature state dependence in one of their examples.

plore misspecification that might arise from arbitrary state dependence in the setting of η_t period-by-period. For this reason we have considered nonseparable specifications of model misspecification with explicit intertemporal tradeoffs.

We achieve computational tractablity in part by our separable specification of an entropy-penality for distortioning q^*. (See the contruction for W in (1).) But this differs from adopting a separable constraint on the date t conditional entropy[9, 10]

$$E\left[\log(q^*_{t+1})q^*_{t+1}|\mathcal{F}_t\right] \le \eta_t.$$

A virtue of the robust control theory approach is that it delivers state dependence in the implied η_t's from a low parameter representation. For instance, we could back-solve η_t from our date zero commitment problem via the formula:

$$\eta_t = r_t - \beta E\left(q^*_{t+1} r_{t+1}|\mathcal{F}_t\right)$$

where $\{r_t\}$ is the date t continuation entropy. However, back-solving for the η's will typically not produce identical decisions and worst case distortions as would emerge from simply exogenously specifying the η's. In the separable constraint specification, the minimization problem for q^*_{t+1} will take account of the fact this choice will alter the probabilities over constraints that will pertain in the future. That will result in different valuation processes and may well lead to a substantively interesting differences between the two approaches.

Nevertheless, this back-solving remains interesting in our investigation because of its links to maximum likelihood estimation and statistical detection. See Anderson, Hansen, and Sargent (2000) for a discussion. Just as a Bayesian explores when a given decision rule is a Bayes rule and evaluates that rule by exploring the implicit prior, we may wish to use the implied $\{\eta_t\}$ process to understand better the probability models that are admitted in robust control problems.[11]

10 Concluding Remarks

In all approaches to robustness and uncertainty aversion, the family of candidate models is *ad hoc*. Savage's single-prior theory and multi-prior generalizations of it are not rich enough to

[9]For sufficiently nice specifications of the state dependence, presumably tractable recursive computation methods can also be developed to solve sparable-constraint models.

[10]By extending the notions of dynamic consistency used by Epstein and Schneider (2001) to include state variables like those that support habit persistence, we suspect that separability in the construction of this constraint will no longer be required. Instead of being specified exogenously, the η_t's will possibly also depend on the same state variables used to capture more familiar forms of time nonseparability. In particular, η_t might depend on past consumptions. Martin Schneider concurred with this guess in private correspondence.

[11]Thus it might illuminate situations in which our continuation entropy approach is not very attractive relative to an approach with an exogenous specification of $\{\eta_t\}$. For instance, if it is optimal to 'zero out' the exposure to risk in some given date, the minimizing agent will chose not to distort beliefs at that date and approximation errors will be allocated in future dates. If the $\{\eta_t\}$ were instead exogenously set to be positive, then multiple beliefs would support the no exposure solution and change substantially the pricing implications.

produce beliefs for alternative hypothetical environments. A virtue of rational expectations is that it delivers one well defined endogenous specification of beliefs, and predicts how beliefs change across environments. Robust control theory does too, although it is not clear that r_0 or η_t should have the status of a policy invariant parameter to be transferred from one environment to another.[12] What is transportable under hypothetical interventions is an important question that can only be addressed with more structure or information from other sources.

Nevertheless, the development of computationally tractable tools for exploring model misspecification and its ramifications for modeling dynamic economies should focus on what are the interesting classes of candidate models for applications. It would impede this endeavor if we were to remove robust control methods from economists' toolkit, since these methods have been designed to be tractable.

References

Anderson, E., L. P. Hansen, and T. Sargent (2000, March). Robustness, detection and the price of risk. Mimeo.

Basar, T. and P. Bernhard (1995). *H_∞-Optimal Control and Related Minimax Design Problems*. Boston: Birkhauser.

Blackwell, D. and M. Girshick (1954). *Thoery of Games and Statistical Decisions*. New York: Wiley.

Chen, Z. and L. G. Epstein (2000). Ambiguity, risk and asset returns in continuous time. unpublished manuscript.

Duffie, D. and L. G. Epstein (1992). Stochastic differential utility. *Econometrica 60*(2), 353–394.

Epstein, L. and M. Schneider (2001, July). Recursive multiple priors. unpublished.

Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics 18*, 141–153.

Hansen, L. P., T. Sargent, and T. Tallarini (1999). Robust permanent income and pricing. *Review of Economic Studies 66*, 873–907.

Hansen, L. P., T. J. Sargent, G. A. Turmuhambetova, and N. Williams (2001). Robustness and uncertainty aversion. unpublished.

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.

[12]But since it can be viewed as a special case that sets $r_0 = 0$, the same qualification applies to rational expectations.

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.

Rustichini, A. (2000, November). unpublished correspondence with the authors.

Wang, T. (2001). Two classes of multi-prior preferences. unpublished.