

Comment

Lars Peter Hansen
University of Chicago and
NBER

1. Introduction

Backus, Routledge and Zin (which I will henceforth refer to as BRZ) have assembled an ambitious catalog and discussion of nonstandard, or *exotic*, specifications of preferences. BRZ include illustrations of how some of these specifications have been used in macroeconomic applications. Collecting the myriad of specifications in a single location is an excellent contribution. It will help to expand the overall accessibility and value of this research.

In my limited remarks, I will not review all of their discussion, but I will develop some themes a bit more and perhaps add a different but complementary perspective on some of the literature. Also, my discussion will feature some contributions not mentioned in the BRZ reader's guide. Most of my discussion will focus on environments in which it is hard or impossible to distinguish seemingly different relaxations of expected utility. While BRZ emphasize more distinctions, I will use some examples to feature similarities across specifications. Much of my discussion will exploit continuous-time limits with Brownian motion information structures to display some revealing limiting cases. In particular, I will draw on contributions not mentioned in the BRZ reader's guide by Duffie and Epstein (1992); Geoffard (1996); Dumas, Uppal, and Wang (2000); Petersen, James, and Dupuis (2000); Anderson, Hansen, and Sargent (2003); and Hansen, Sargent, Turmuhambetova, and Williams (2004) along with some of the papers cited by BRZ.

As a precursor to understanding the new implications of exotic preferences, we explore how seemingly different motivations for altering preferences give rise to similar implications and in some circumstances the same implications. BRZ have separate sections entitled *time* (Section 2), *time and risk* (Section 4), *risk sensitive and robust control* (Section

5), and *ambiguity* (Section 6). In what follows, I will review some existing characterizations in the literature to display a tighter connection than what might be evident from reading their paper.

2. Endogenous Discounting

I begin with a continuous-time version of the discussion in the BRZ treatment of *time* (Section 2 of their paper). An important relaxation of discounted utility is the recursive formulation of preferences suggested by Koopmans (1960), Uzawa (1968), and others. These are preferences that allow for endogenous discounting. A convenient generalization of these preferences is one in which the discount rate is a choice variable subject to a utility penalty, as in the *variational utility* specification of Geoffard (1996).

Consider preferences for consumption defined over an interval of time $[0, T]$ with undiscounted continuation value U_t that satisfies:

$$\begin{aligned}\lambda_t U_t &= E_t \int_t^T \lambda_s F(c_s, v_s) ds \\ \lambda_t &= \exp \left(\int_0^t -v_\tau d\tau \right)\end{aligned}\tag{1}$$

where $\{c_t : 0 \leq t \leq T\}$ is an admissible consumption process and $\{v_t : 0 \leq t \leq T\}$ is an admissible subjective discount rate process.¹ Then λ_t is a discount factor constructed from current and past discount rates. The notation E_t is used to denote the expectation operator conditioned on date t information. Equation (1) determines the continuation values for a consumption profile for each point in time. In particular, the date zero utility function is given by:

$$U_0 = E_0 \int_0^T \lambda_s F(c_s, v_s) ds$$

The function F gives the instantaneous contribution to utility, and it can depend on the subjective rate of discount v_s for reasons that will become clear.

So far we have specified the discounting in a flexible way, but stipulating the subjective discount rates must still be determined.² To convert this decision problem into an endogenous discount factor model, we follow Geoffard (1996) by determining the discount rate via minimization. This gives rise to a nondegenerate solution because of our

choice to enter v as an argument in the function F . To support this minimization, the function $F(c, v)$ is presumed to be convex in v . Given the recursive structure to these preferences, v solves the continuous-time Bellman equation:

$$V(c_t, U_t) \doteq \inf_v [F(c_t, v) - vU_t] \quad (2)$$

The first-order conditions for minimizing v are:

$$F_v(c_t, v_t) = U_t$$

which implicitly defines the discount rate v_t as a function of the current consumption c_t and the current continuation value U_t .

This minimization also implies a forward utility recursion in U_t by specifying its drift:

$$\lim_{\varepsilon \downarrow 0} \frac{E_t U_{t+\varepsilon} - U_t}{\varepsilon} = -V(c_t, U_t)$$

This limit depicts a Koopmans (1960)–style aggregator in continuous-time with uncertainty. Koopmans (1960) defined an implied discount factor via a differentiation. The analogous implicit discount rate is given by the derivative:

$$v = -V_U(c, U)$$

consistent with representation (1).

So far we have seen how a minimum discount rate formulation implies an aggregator of the type suggested by Koopmans (1960) and others. As emphasized by Geoffard (1996), we may also go in the other direction. Given a specification for V , the drift for the continuation value, we may construct a Geoffard (1996)–style aggregator. This is accomplished by building a function F from the function V . The construction (2) of V formally is the Legendre transform of F . This transform has an inverse given by the algorithm:

$$F(c, v) = \sup_U [V(c, U) + vU] \quad (3)$$

Example 2.1 The implied discount rate is constant and equal to δ when:

$$V(c, U) = u(c) - \delta U$$

Taking the inverse Legendre transform, it follows that:

$$\begin{aligned}
 F(c, v) &= \sup_U [u(c) - \delta U + vU] \\
 &= \begin{cases} u(c) & \text{if } v = \delta \\ +\infty & \text{if } v \neq \delta \end{cases}
 \end{aligned}$$

This specification of V and F gives rise to the familiar discounted utility model.

Of course, the treatment of *exotic* preferences leads us to explore other specifications outside the confines of this example. These include preferences for which v is no longer constant.

In economies with multiple consumers, a convenient device to characterize and solve for equilibria is to compute the solutions to resource allocation problems with a social objective given by the weighted sum of the individual utility functions (Negishi, 1960). As reviewed by BRZ, Lucas and Stokey (1984) develop and apply an intertemporal counterpart to this device to study economies in which consumers have recursive utility. For a continuous time specification, Dumas, Uppal, and Wang (2000) use Geoffard's formulation of preferences to characterize efficient resource allocations. This approach also uses Negishi/Pareto weights and discount rate minimization. Specifically Dumas, Uppal, and Wang (2000) use a social objective:

$$\begin{aligned}
 \inf_{\{v_t^i: \tau \geq t\}} \sum_i E_t \int_t^T \lambda_s^i F^i(c_s^i, v_s^i) ds \\
 \frac{d\lambda_t^i}{dt} = -v_t^i \lambda_t^i
 \end{aligned} \tag{4}$$

where the Negishi weights are the date zero initial conditions for λ_0^i and i denotes individuals.

Thus far, we have produced two ways to represent endogenous discount factor formulations of preferences. BRZ study the Koopmans (1960) specification in which $V(c, u)$ is specified and a discount rate is defined as $-V_U(c, U)$. In the Geoffard (1996) characterization, $V(c, U)$ is the outcome of a problem in which discounted utility is minimized by choice of a discount rate process. The resulting function is concave in U . As we will see, however, the case in which V is convex in U is of particular interest to us. An analogous development to that given by Geoffard (1996) applies in which discounted utility is maximized by choice of the discount rate process instead of minimized.

3. Risk Adjustments in Continuation Values

Consider next a specification of preferences due to Kreps and Porteus (1978) and Epstein and Zin (1989). (BRZ refer to these as *Kreps–Porteus preferences* but certainly Epstein and Zin played a prominent role in demonstrating their value.) In discrete time, these preferences can be depicted recursively using a recursion with a risk-adjustment to the continuation value of the form:

$$U_t^* = u(c_t) + \beta h^{-1} E_t h(U_{t+1}^*) \quad (5)$$

As proposed by Kreps and Porteus (1978), the function h is increasing and is used to relax the assumption that compound intertemporal lotteries for utility can be reduced in a simple manner. When the function h is concave, it enhances risk aversion without altering intertemporal substitution (see Epstein and Zin, 1989).

Again it is convenient to explore a continuous-time counterpart. To formulate such a limit, scale the current period contribution by ε , where ε is the length of the time interval between observations, and parameterize the discount factor β as $\exp(-\delta\varepsilon)$, where δ is the instantaneous subjective rate of discount. The local version of the risk adjustment is:

$$\lim_{\varepsilon \downarrow 0} \frac{E_t h(U_{t+\varepsilon}^*) - h(U_t^*)}{\varepsilon} = -h'(U_t^*) [u(c_t) - \delta U_t^*] \quad (6)$$

The lefthand side can be defined for a Brownian motion information structure and for some other information structures that include jumps.

Under a Brownian motion information structure, the local evolution for the continuation value can be depicted as:

$$dU_t^* = \mu_t^* dt + \sigma_t^* \cdot dB_t \quad (7)$$

where $\{B_t\}$ is multivariate standard Brownian motion. Thus, μ_t^* is the local mean of the continuation value and $|\sigma_t^*|^2$ is the local variance:

$$\mu_t^* = \lim_{\varepsilon \downarrow 0} \frac{E_t U_{t+\varepsilon}^* - U_t^*}{\varepsilon}$$

$$|\sigma_t^*|^2 = \lim_{\varepsilon \downarrow 0} \frac{E_t (U_{t+\varepsilon}^* - U_t^*)^2}{\varepsilon}$$

By Ito's Lemma, we may compute the local mean of $h(U_t^*)$:

$$\frac{E_t h(U_{t+\varepsilon}^*) - h(U_t^*)}{\varepsilon} = h'(U_t^*)\mu_t^* + \frac{1}{2}h''(U_t^*)|\sigma_t^*|^2$$

Substituting this formula into the lefthand side of equation (6) and solving for μ_t^* gives:

$$\mu_t^* = \delta U_t^* - u(c_t) - \frac{h''(U_t^*)}{2h'(U_t^*)}|\sigma_t^*|^2 \quad (8)$$

Notice that the risk-adjustment to the value function adds a variance contribution to the continuation value recursion scaled by what Duffie and Epstein (1992) refer to as the *variance multiplier*, given by:

$$\frac{h''(U_t^*)}{h'(U_t^*)}$$

When h is strictly increasing and concave, this multiplier is negative. The use of h as a risk adjustment of the continuation value gives rise to concern about variation in the continuation value. Both the local mean and the local variance are present in this recursion.

As Duffie and Epstein (1992) emphasize, we can transform the utility index and eliminate the explicit variance contribution. Applying such a transformation gives an explicit link between the Kreps and Porteus (1978) specification and the Koopmans (1960) specification. To demonstrate this, transform the continuation value via $U_t = h(U_t^*)$. This results in the formula:

$$\lim_{\varepsilon \downarrow 0} \frac{E_t U_{t+\varepsilon} - U_t}{\varepsilon} = -V(c_t, U_t)$$

where

$$V(c, U) = h'[h^{-1}(U)][u(c) - \delta h^{-1}(U)]$$

The Geoffard (1996) specification with discount rate minimization can be deduced by solving for the inverse Legendre transform in equation (3). The implied endogenous discount rate is:

$$-V_U(c, U) = \delta - \frac{h''[h^{-1}(U)]}{h'[h^{-1}(U)]}[u(c) - \delta h^{-1}(U)]$$

Consider two examples. The first has been used extensively in the literature linking asset prices and macroeconomics aggregates including consumption.

Example 3.1 Consider the case in which

$$u(c) = \frac{c^{1-\varrho}}{1-\varrho} \quad \text{and} \quad h(U^*) = \frac{[(1-\varrho)U^*]^{(1-\gamma)/(1-\varrho)}}{1-\gamma}$$

where $\varrho > 0$ and $\gamma > 0$. We assume that $\varrho \neq 1$ and $\gamma \neq 1$ because the complementary cases require some special treatment. This specification is equivalent to the specification given in equations (9) and (10) of BRZ.³ Then:

$$V(c, U) = [(1-\gamma)U]^{(\varrho-\gamma)/(1-\gamma)} \left(\frac{c^{1-\varrho}}{1-\varrho} \right) - \delta \left(\frac{1-\gamma}{1-\varrho} \right) U$$

with implied endogenous discount rate:

$$v = \frac{1-\gamma}{1-\varrho} \delta + \frac{(\gamma-\varrho)}{(1-\varrho)} \left[\frac{u(c)}{h^{-1}(U)} \right]$$

Notice that the implied endogenous discount rate simplifies, as it should, to be δ when $\varrho = \gamma$. The dependent component of the discount rate depends on the discrepancy between ϱ and γ and on the ratio of the current period utility to the continuation value without the risk adjustment:

$$U^* = h^{-1}(U)$$

At the end of Section 2, we posed an efficient resource allocation problem (4) with heterogeneous consumers. In the heterogeneous consumer economy with common preferences of the form given in Example 3.1, the consumption allocation rules as a function of aggregate consumption are invariant over time. The homogeneity discussed in Duffie and Epstein (1992) and by BRZ implies that the ratio of current period utility to the continuation value will be the same for all consumers, implying in turn that the endogenous discount rates will be also. With preference heterogeneity, this ceases to be true, as illustrated by Dumas, Uppal, and Wang (2000).

We will use the next example to relate to the literature on robustness in decisionmaking. It has been used by Tallarini (1998) in the study of business cycles and by Anderson (2004) to study resource allocation with heterogeneous consumers.

Example 3.2 Consider the case in which $h(U^) = -\theta \exp(-U^*/\theta)$ for $\theta > 0$. Notice that the transformed continuation utility is negative. A simple calculation results in:*

$$V(c, U) = -\frac{U}{\theta} \left[u(c) + \delta \theta \log \left(-\frac{U}{\theta} \right) \right]$$

which is convex in U . The maximizing v of the Legendre transform (2) is:

$$v = \delta + \frac{1}{\theta} \left[u(c) + \delta \theta \log \left(-\frac{U}{\theta} \right) \right]$$

and the minimizing U of the inverse Legendre transform (3) is:

$$U = -\theta \exp \left[\frac{\theta v - \theta \delta - u(c)}{\delta \theta} \right]$$

Consequently:

$$F(c, v) = -\delta \theta \exp \left[\frac{\theta v - \theta \delta - u(c)}{\delta \theta} \right]$$

which is concave in v .

So far, we have focused on what BRZ call Kreps–Porteus preferences. BRZ also discuss what they call Epstein–Zin preferences, which are dynamic recursive extensions to specifications of Chew (1983) and Dekel (1986). Duffie and Epstein (1992) show, however, how to construct a corresponding variance multiplier for versions of these preferences that are sufficiently smooth and how to construct a corresponding risk-adjustment function h for Brownian motion information structures (see page 365 of Duffie and Epstein, 1992).

This equivalence does not extend to all of the recursive preference structures described by BRZ. This analysis has not included, for instance, dynamic versions of preferences that display first-order risk aversion.⁴ BRZ discuss such preferences and some of their interesting implications.

Let me review what has been established so far. By taking a continuous-time limit for a Brownian motion information structure, a risk-adjustment in the continuation value for a consumption profile is equivalent to an endogenous discounting formulation. We can view this endogenous discounting as a continuous-time version of a Koopmans (1960)–style recursion or as a specification in which discount rates are the solution to an optimization problem, as in Geoffard (1996). These three different starting points can be used to motivate the same set of preferences. Thus, we produced examples in which some of the preference specifications in Sections 2 and 4 of BRZ are formally the same.

Next, we consider a fourth specification.

4. Robustness and Entropy

Geoffard (1996) motivates discount rate minimization as follows:

[T]he future evolution of relevant variables (sales volumes, asset default rates or prepayment rates, *etc.*) is very important to the valuation of a firm's debt. A probability distribution on the future of these variables may be difficult to define. Instead, it may be more intuitive to assume that these variables remain within some confidence interval, and to define the value of the debt as the value in the worst case, *i.e.* when the evolution of the relevant state variables is systematically adverse.

It is not obvious that Geoffard's formalization is designed for a robustness adjustment of this type. In what follows a conservative assessment made by exploring alternative probability structures instead leads to a formulation where the discounted utility is maximized by choice of discount rates and not minimized because the implied $V(c, U)$ is convex in U . In this section we will exploit a well-known close relationship between *risk sensitivity* and a particular form of *robustness* from control theory, starting with Jacobson (1973). A discussion of the linear-quadratic version of risk-sensitive and robust control theory is featured in Section 5 of BRZ. The close link is present in much more general circumstances, as I now illustrate.

Instead of recursion (5), consider a specification in which beliefs are distorted subject to penalization:

$$U_t^* = \min_{q_{t+1} \geq 0, E_t q_{t+1} = 1} u(c_t) + \beta E_t(U_{t+1}^* q_{t+1}) + \beta \theta E_t[(\log q_{t+1}) q_{t+1}] \quad (9)$$

The random variable q_{t+1} distorts the conditional probability distribution for date $t + 1$ events conditioned on date t information. We have added a penalization term to limit the severity of the probability distortion. This penalization is based on a discrepancy measure between the implied probability distributions called conditional relative entropy. Minimizing with respect to q_{t+1} in this specification produces a version of recursion (5), with h given by the risk-sensitive specification of Example 3.2. It gives rise to the exponential tilting because the penalized worst-case q_{t+1} is:

$$q_{t+1} \propto \exp\left(-\frac{U_{t+1}^*}{\theta}\right)$$

Probabilities are distorted less when the continuation value is high and more when this value is low. By making the θ large, the solution to this

problem approximates that of the recursion of the standard form of time-separable preferences. Given this dual interpretation, robustness can look like risk aversion in decisionmaking and in prices that clear security markets. This dual interpretation is applicable in discrete and continuous time. For a continuous time analysis, see Hansen, Sargent, Turmuhambetova, and Williams (2004) and Skiadas (2003).

Preferences of this sort are supported by worst-case distributions. Blackwell and Girshick (1954) organize statistical theory around the theory of two-player zero-sum games. This framework can be applied in this environment as well. In a decision problem, we would be led to solve a max-min problem. Whenever we can exchange the order of minimization and maximization, we can produce a worst-case distribution for the underlying shocks under which the action is obtained by a simple maximization. Thus, we can produce ex post a shock specification under which the decision process is optimal and solves a standard dynamic programming problem. It is common in Bayesian decision theory to ask what prior justifies a particular rule as being optimal. We use the same logic to produce a (penalized) worst-case specification of shocks that justifies a robust decision rule as being optimal against a correctly specified model.

This poses an interesting challenge to a rational expectations econometrician studying a representative agent model. If the worst-case model of shock evolution is statistically close to that of the original model, then an econometrician will have difficulty distinguishing *exotic preferences* from a possibly more complex specification of shock evolution. See Anderson, Hansen, and Sargent (2003) for a formal discussion of the link between statistical discrimination and robustness and Hansen, Sargent, Turmuhambetova, and Williams (2004) for a discussion and characterization of the implied worst-case models for a Brownian motion information structure. In the case of a decision problem with a diffusion specification for the state evolution, the worst-case model replaces the Brownian motion shocks with a Brownian motion distorted by a nonzero drift.

In the case of Brownian motion information structures, Maenhout (2004) has shown the robust interpretation for a more general class of recursive utility models by allowing for a more general specification of the penalization. Following Maenhout (2004), we allow θ to depend on the continuation value U_t^* .

In discrete time, we distorted probabilities using a positive random variable q_{t+1} with conditional expectation equal to unity. The product of such random variables:

$$z_{t+1} = \prod_{j=1}^{t+1} q_j$$

is a discrete time martingale. In continuous time, we use nonnegative martingales with unit expectations to depict probability distortions. For a Brownian motion information structure, the local evolution of a nonnegative martingale can be represented as:

$$dz_t = z_t g_t \cdot dW_t$$

where g_t dictates how the martingale increment is related to the increment in the multivariate Brownian motion $\{W_t : t \geq 0\}$. In continuous time, the counterpart to $E_t(q_{t+1} \log q_{t+1})$ is the quadratic penalty $|g_t|^2/2$, and our minimization will entail a choice of the random vector g_t .

In accordance with Ito's formula, the local mean of the distorted expectation of the continuation value process $\{U_t^* : t \geq 0\}$ is:

$$\lim_{\varepsilon \downarrow 0} \frac{E_t z_{t+\varepsilon} U_{t+\varepsilon}^* - z_t U_t^*}{\varepsilon} = z_t \mu_t^* + z_t \sigma_t^* g_t$$

where the continuation value process evolves according to equation (7). The continuous-time counterpart to equation (9) is:

$$z_t \mu_t^* = \min_{g_t} -z_t \sigma_t^* g_t - z_t u(c_t) + z_t \delta U_t^* - z_t \theta(U_t^*) \frac{|g_t|^2}{2}$$

with the minimizing value of g_t given by:

$$g_t = -\frac{\sigma_t^*}{\theta(U_t^*)}$$

Substituting for this choice of g_t , the local mean for the continuation value must satisfy:

$$\mu_t^* = -u(c_t) + \delta U_t^* + \frac{|\sigma_t^*|^2}{2\theta(U_t^*)}$$

(provided of course that z_t is not zero). By setting θ to be:

$$\theta(U^*) = -\frac{h'(U^*)}{h''(U^*)}$$

we reproduce equation (8) and hence obtain the more general link among utility recursions for h increasing and concave. This link,

however, has been established only for a continuous-time economy with a Brownian motion information structure for a general specification of h .

The penalization approach can nest other specifications not included by the utility recursions I discussed in Sections 2 and 3. For instance, the concern about misspecification might be concentrated on a proper subset of the shock processes (the Brownian motions).

To summarize, we have now added a concern about model specification to our list of *exotic preferences* with comparable implications when information is approximated by a Brownian motion information structure. When there is a well-defined worst-case model, an econometrician might have trouble distinguishing these preferences from a specification with a more complex but statistically similar evolution for the underlying economic shocks.

5. Uncertainty Aversion

The preferences built in Section 4 were constructed using a penalty based on conditional relative entropy. Complementary axiomatic treatments of this penalty approach to preferences have been given by Wang (2003) and Maccheroni, Marinacci, and Rustichini (2004).

Formulation (9) used θ as a penalty parameter, but θ can also be the Lagrange multiplier on an intertemporal constraint (see Petersen, James, and Dupuis, 2000, and Hansen, Sargent, Turmuhambetova, and Williams, 2004). This interpretation of θ as a Lagrange multiplier links our previous formulation of robustness to decision making when an extensive family of probability models are explored subject to an intertemporal entropy constraint. While the implied preferences differ, the interpretation of θ as a Lagrange multiplier gives a connection between the decision rules from the robust decision problem described at the outset of Section 4 and the multiple priors model discussed in Section 6 of BRZ. Thus, we have added another possible interpretation to the risk-sensitive recursive utility model. Although the Lagrange multiplier interpretation is deduced from a date zero vantage point, Hansen, Sargent, Turmuhambetova, and Williams (2004) describe multiple ways in which such preferences can look recursive.

Of course, there are a variety of other ways in which multiple models can be introduced into a decision problem. BRZ explore some aspects of dynamic consistency as it relates to decision problems with multiple probability models. A clear statement of this issue and its

ramifications requires much more than the limited space BRZ had to address it. As a consequence, I found this component of the paper less illuminating than other components.

A treatment of dynamic consistency with multiple probability models either from the vantage point of robustness or ambiguity is made most interesting by the explicit study of environments in which learning about a parameter or a hidden state through signals is featured. Control problems are forward-looking and are commonly solved using a backward induction method such as dynamic programming. Predicting unknown states or estimating parameters is inherently backward-looking. It uses historical data to make a current period prediction or estimate. In contrast to dynamic programming, recursive prediction iterates going forward. This difference between control and prediction is the source of tension when multiple probability models are entertained. Recursive formulations often ask that you back away from the search for a single coherent worst-case probability model over observed signals and hidden states or parameters. The connection to Bayesian decision theory that I mentioned previously is often broken. In my view, a pedagogically useful treatment of this issue has yet to be written, but it requires a separate paper.

6. Conclusion

We have shown how divergent motivations for generalizing preferences sometimes end up with the same implications. So what? There are at least three reasons I can think of why an economic researcher should be interested in these alternative interpretations. One reason is to understand how we might calibrate or estimate the new preference parameters. The different motivations might lead us to think differently about what is a reasonable parameter setting. For instance, what might appear to be endogenous discounting could instead reflect an aversion to risk when a decision maker cares about the intertemporal composition of risk. What might look like an extreme amount of risk aversion could instead reflect the desire of the decision maker to accommodate model misspecification.

Second, we should understand better the new testable implications that might emerge as a result of our exploring nonstandard preferences. Under what auxiliary assumptions are there interesting testable implications? My remarks point to some situations when testing will be challenging or fruitless.

Finally, we should understand better when preference parameters can be transported from one environment to another. This understanding is at least implicitly required when we explore hypothetical changes in macroeconomic policies.

It would be nice to see a follow-up paper that treated systematically (1) the best sources of information for the new parameters, (2) the observable implications, and (3) the policy consequences.

Notes

Conversations with Jose Mazoy, Monika Piazzesi, and Grace Tsiang were valuable in the preparation of these remarks.

1. We may define formally the notion of *admissible* by restricting the consumption and discount rate processes to be progressively measurable given a prespecified filtration.
2. Geoffard (1996) does not include uncertainty in his analysis, but as Dumas, Uppal, and Wang (2000) argue, this is a straightforward extension.
3. This equivalence follows by letting $\rho = 1 - \varrho$ and $\alpha = 1 - \gamma$ and transforming the utility index.
4. See Duffie and Epstein (1992), page 361, for a more complete discussion about what is excluded under the Brownian information structure by their variance multiplier formulation.

References

- Anderson, E. (2004). The dynamics of risk-sensitive allocations. University of North Carolina. Unpublished Manuscript.
- Anderson, E., L. Hansen, and T. Sargent. (2003). A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association* 1(1):68–123.
- Blackwell, D., and M. A. Girshick. (1954). *Theory of Games and Statistical Decisions*. New York: Wiley Publications in Statistics.
- Chew, S. H. (1983). A generalization of the quasi-linear mean with applications to the measurement of inequality and decision theory resolving the Allais paradox. *Econometrica* 51:1065–1092.
- Dekel, E. (1986). An axiomatic characterization of preference under uncertainty: Weakening the independence axiom. *Journal of Economic Theory* 40:304–318.
- Duffie, D., and L. G. Epstein. (1992). Stochastic differential utility. *Econometrica* 60(2):353–394.
- Dumas, B., R. Uppal, and T. Wang. (2000). *Journal of Economic Theory* 93:240–259.
- 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.

Geoffard, P. Y. (1996). Discounting and optimizing: Capital accumulation problems as variational minmax problems. *Journal of Economic Theory* 69:53–70.

Hansen, L. P., T. J. Sargent, G. A. Turmuhambetova, and N. Williams. (2004). Robust control, min-max expected utility, and model misspecification. University of Chicago. Unpublished Manuscript.

Jacobson, D. H. (1973). Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games. *IEEE Transactions for Automatic Control* AC-18:1124–1131.

Koopmans, T. J. (1960). Stationary ordinal utility and impatience. *Econometrica* 28:287–309.

Kreps, D. M., and E. L. Porteus. (1978). Temporal resolution of uncertainty and dynamic choice. *Econometrica* 46:185–200.

Lucas, R. E., and N. L. Stokey. (1984). Optimal growth with many consumers. *Journal of Economic Theory* 32:139–171.

Maccheroni, F., M. Marinacci, and A. Rustichini. (2004). Variational representation of preferences under ambiguity. University of Minnesota. Unpublished Manuscript.

Maenhout, P. J. (2004). Robust portfolio rules and asset pricing. *Review of Financial Studies* 17:951–983.

Negishi, T. (1960). Welfare economics and existence of an equilibrium for a competitive economy. *Metroeconomica* 12:92–97.

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.

Skiadas, C. (2003). Robust control and recursive utility. *Finance and Stochastics* 7:475–489.

Tallarini, T. (1998). Risk sensitive business cycles. *Journal of Monetary Economics* 43:507–532.

Uzawa, H. (1968). Time preference, the consumption function, and optimum asset holdings. In *Value, Capital, and Growth: Papers in Honor of Sir John Hicks*, J. N. Wolfe (ed.). Edinburgh: Edinburgh University.

Wang, T. (2003). A class of multi-prior preferences. University of British Columbia. Unpublished Manuscript.