Notes on Robust Portfolio Choice*

Parag A. Pathak†

Preliminary; Comments Welcome

First version: 03/03/02 Current version: 12/07/02

Abstract. There have been many recent attempts to extend the static max-min expected utility theory of Gilboa and Schmeidler to dynamic environments. The purpose of this note is to re-direct attention in portfolio choice to frameworks which are axiomatized and motivated by a corresponding theory of choice. In particular, we argue that papers including Maenhout (2001), Uppal and Wang (2003), and Liu, Pan, and Wang (2002) employ a transformation that is poorly motivated and breaks the link to a foundation based on Gilboa and Schmeidler. We show, instead, how the analytical results of these papers can be derived in the *recursive multiple priors* framework that has been axiomatized and is clearly linked to Gilboa-Schmeidler. While the intuition for the results are similar, the comparative statics are significantly different. Finally, we offer some broader thoughts about whether Gilboa and Schmeidler is a convincing nonexpected utility alternative for portfolio choice.

Key Words: robust control, portfolio choice

* This paper is extracted from my senior thesis entitled: "Portfolio Choice, Uncertainty Aversion, and Limited Diversification" submitted to the committee on Applied Mathematics at Harvard University. I am particularly grateful to G. Chamberlain for supervising my senior thesis and teaching me how to do research. L. Epstein, T. Knox and T. Wang provided detailed comments and helpful discussion. T. Anderson, J. Campbell, G. Chacko, D. Fudenberg, T. Knox, D. Laibson, P. Maenhout, J. Pan, C. Skiadas, R. Uppal, L. Viceira, and N. Williams also provided helpful comments. Some of this work was done with the financial support of the Chateaubriand Fellowship.

† Harvard University, e-mail: ppathak@fas.harvard.edu

1. INTRODUCTION

There have been many recent attempts to extend the static maxmin expected utility theory of Gilboa and Schmeidler (1989) to a dynamic environment useful for econometrics, finance, and macroeconomics. Unlike subjective expected utility, the extension of static non-expected utility preferences to intertemporal lotteries is a nontrivial undertaking.[1] The main purpose of this paper is to critically assess recent contributions to *robust* portfolio choice theory, building on and modifying the framework proposed by Lars Hansen, Thomas Sargent, and a set of authors. Our argument is that the utility specifications employed in a set of papers—namely Maenhout (2001), Uppal and Wang (2003), and Liu, Pan, and Wang (2002)—have lost their link to the motivating decision theory of Gilboa and Schmeidler.

The argument centers on a transformation, initially proposed by Maenhout (2001). Expressed in terms of the framework of Hansen, Sargent, Turmuhambetova, and Williams (2002), Maenhout transforms a constant Lagrange multiplier into a function of the value function for analytical tractability. The Hansen-Sargent framework, however, has a constant Lagrange multiplier regulating the class of models the decision maker considers in order to motivate their preferences by Gilboa-Schmeidler. After making this explicit, we probe deeper into the implications of Maenhout's transformation and show that all existing applications relying on Maenhout's transformation can be expressed in terms of the *recursive multiple priors* framework for intertemporal Gilboa-Schmeidler axiomatized in Epstein and Schneider (2002b). We obtain analytical solutions with similar intuitions, but different comparative statics. In the process of establishing this point, we show a connection to the *robust control* framework. Lastly, we present some comments about whether Gilboa-Schmeidler is a convincing non-expected utility foundation for portfolio choice.

Gilboa and Schmeidler's theory provides a tractable way of modeling decision makers who display an aversion to uncertainty or am-

[1] See, for example, the class of recursive, but not necessarily expected utility preferences in Epstein and Zin (1989).

biguity and may desire a decision rule that is robust to model misspecification. There are many approaches to intertemporal preferences which are consistent with Gilboa-Schmeidler. Three different approaches have been proposed by (1) Larry Epstein and a set of co-authors, (2) Lars P. Hansen, Thomas J. Sargent, and a set of co-authors and (3) Gary Chamberlain and Thomas A. Knox.[2]

The recursive multiple priors model has been introduced and studied by Epstein and a set of co-authors.[3] Another framework, introduced by Hansen and Sargent and a set of co-authors,[4] is based on robust control theory and motivated by Gilboa and Schmeidler. Their main paper introduces two classes of preferences[5] and provides easy-to-use methods to find optimal robust decision rules using tools from dynamic programming. Lastly, Chamberlain (2000) presents a framework and numerical algorithm, and Knox (2002) axiomatizes and generalizes it. These three frameworks all have similar implications, but they have subtle and important differences. Early applications of all three frameworks have studied the canonical problem of portfolio choice.

Maenhout (2001) adopts and modifies the Hansen-Sargent framework to study Merton's (1969) optimal portfolio choice problem. In the next section, we briefly review the Hansen, Sargent, Tur-

[2]Knox (2002), unlike the other approaches, does not constitute a joint project. Other models include Klibanoff (1995), Siniscalchi (2001), and the framework provided by Wang (2001) combined with Wang (2002) which falls in the Hansen-Sargent class.

[3]Epstein and Wang (1994) first introduced a nonaxiomatic version of recursive multiple priors in discrete-time, Chen and Epstein (2002) present a continuous-time version of multi-prior utility, Epstein and Schneider (2002b) present an axiomatization of the recursive priors model, and Epstein and Schneider (2002a) present an application with learning and multiple priors. We will hereafter refer to the model as recursive multiple priors.

[4]See Hansen and Sargent (2001b), Hansen and Sargent (2001c), Hansen and Sargent (2001d), Anderson, Hansen, and Sargent (2002), Hansen, Sargent, Turmuhambetova, and Williams (2002), and the manuscript Hansen and Sargent (2001a). We will hereafter refer to this as the robust control model and focus only on the multiplier preferences presented in Hansen, Sargent, Turmuhambetova, and Williams (2002).

[5]An axiomatization of one of these preferences - the "multiplier" preferences can be constructed by combining Wang (2001) and Wang (2002).

muhambetova, and Williams (2002) (hereafter, known as HSTW)[6] robust-control formulation of Gilboa and Schmeidler (1989). Section 3 discusses Maenhout's (2001) adaptation of Merton's (1969) portfolio choice problem for a robust investor and both Uppal and Wang (2003) and Liu, Pan, and Wang (2002) who extend this foundation. We aim to clarify why the transformation breaks a link to the foundation of Gilboa and Schmeidler. Section 4 describes the alternative recursive multiple priors model and obtains closed-form rules for the portfolio choice problem. In section 5, we show how to obtain analytical solutions for existing applications employing Maenhout's (2001) transformation using recursive multiple priors. In section 6, we offer some general remarks about the direction of this literature. The last section concludes.

2. OVERVIEW OF ROBUST CONTROL

Motivated by Knight (1921) and Ellsberg (1961), Gilboa and Schmeidler (1989) present axioms of choice and a representation theorem for an individual with preferences over a set of subjective priors. The theory is known as max-min expected utility because the agent evaluates lotteries by calculating the expected utility for each distribution $Q \in \mathcal{Q}$ and makes a comparision according to the minimum value of this index. Gilboa and Schmeidler's axiomatization is presented for lotteries defined in the Anscombe and Aumann (1963) domain, but has been recently adapted to Savage acts by Casadesus-Masanell, Klibanoff, and Ozdenoren (2000).

Define the states of the world by Ω, with algebra of subsets Σ and denote the set of outcomes or prizes by X where $\mathcal{Y} = \Delta(X)$, the set of finite-support probability distributions on X. The set of *lottery-acts* $\mathcal{L}$ is defined as the set of Σ-measurable finite step functions from Ω to $\mathcal{Y}$. We are interested in comparing lottery-acts of the form $f \in \mathcal{L}$ where $f : \Omega \to \mathcal{Y}$, over which individuals have a given preference relation $\succeq$. Gilboa-Schmeidler's representation result relies on axioms which are typical in decision theory with the

[6] While there are some differences across the Hansen-Sargent work, we focus on Hansen, Sargent, Turmuhambetova, and Williams (2002) because it places the most emphasis on a decision-theoretic approach motivated by Gilboa-Schmeidler.

exception of a weakened independence axiom.[7] Their result states that for some closed and convex set of distributions $\mathcal{Q}$, there exists an affine function $u : \mathcal{Y} \to \mathbb{R}$ such that for all $f, g \in \mathcal{L}$,

$$f \succeq g \quad \Leftrightarrow \quad \inf_{Q\in\mathcal{Q}} \int_\Omega u(f(\omega))dQ(w) \succeq \inf_{Q\in\mathcal{Q}} \int_\Omega u(g(\omega))dQ(w). \quad (1)$$

The challenge is taking this static theory and suitably adapting it to intertemporal lotteries and placing a reasonable restriction on the set $\mathcal{Q}$.

In HSTW, an intertemporal version of the Gilboa-Schmeidler decision theory is operationalized through the techniques of robust control theory. HSTW posit that Gilboa and Schmeidler's decision theory implies preferences over uncertainty of the form:

$$\inf_{Q\in\mathcal{Q}} \mathbb{E}_Q \left[\int_0^\infty e^{-\delta t} U(c_t, x_t) dt \right], \quad (2)$$

where $U(\cdot)$ is the utility function, c_t is consumption, x_t is the state variable, and δ is the subjective discount factor. We can re-write the utility process in terms of c_t alone once we have specified the dynamics of x_t. This will constitute an objective which is an infinite-horizon time-additive specification of Gilboa-Schmeidler. Specifying the dynamics of x_t goes hand-in-hand with delineating the set of distributions $\mathcal{Q}$. In robust control theory, the set $\mathcal{Q}$ is by delineating through perturbations to the underlying objective reference model.

The benchmark problem examined by HSTW specifies a continuous-time diffusion for the state evolution governed by $\{Z_t : t \geq 0\}$, a d-dimensional Brownian motion on an underlying probability space $(\Omega, \mathcal{F}, P)$. The control problem of the decision maker is:

$$\sup_{c\in C} \mathbb{E} \left[\int_0^\infty e^{-\delta t} U(c_t, x_t) dt \right] \quad (3)$$

subject to state evolution

$$dx_t = \mu(c_t, x_t)dt + \sigma(c_t, x_t)dZ_t, \quad (4)$$

with given initial condition. HSTW specify the set $\mathcal{Q}$ as the set of absolutely continuous models with respect to the reference distribution.

[7]See Gilboa and Schmeidler (1989) for specific details.

This assumption is made for two reasons: (1) to ensure that perturbations are difficult to detect with finite amounts of data and (2) to employ information-based measures of distance, such as entropy. Absolute continuity requires that measures agree on zero-probability events, and in the context of a continuous-time diffusion, Girsanov's Theorem implies that absolute continuity reduces to altering the drift of the underlying model. In other words, the perturbations specify that outside a set of Q-measure zero, the distribution Q has the same form as (4) with the Brownian increment dZ_t replaced by $g_t dt + d\hat{Z}_t$, where g_t is a progressively measurable stochastic process with the same dimension as Z_t and $d\hat{Z}_t$ is a Brownian increment under Q. The set of alternative measures $\mathcal{Q}$ on $(\Omega, \mathcal{F})$ refers to different specifications of the stochastic process $\{g_t : t \geq 0\}$. For a particular $Q \in \mathcal{Q}$, the distorted system can be expressed as:

$$dx_t = \mu(c_t, x_t)dt + \sigma(c_t, x_t)(g_t dt + d\hat{Z}_t).$$

The set of admissable distributions is governed by a constraint on the relative entropy or the Kullback-Leibler information divergence between a candidate distribution $Q \in \mathcal{Q}$ and the reference distribution P. Informally, this measure is similar to a log likelihood ratio, though it is not a true metric because it is not symmetric and does not satisfy the triangle inequality. There is no firm foundation for the choice of the entropy discrepancy; other alternatives for measuring discrepancy exist, but will likely lead to the same implications. In the context of our continuous-time diffusion model, HSTW have defined relative entropy as:

$$\mathcal{R}(Q \parallel P) \equiv \delta \int_0^\infty e^{-\delta t} \mathbb{E}_Q[\log q_t] dt,$$

where $q_t = \mathbb{E}[q|\mathcal{F}_t]$ and $\frac{dQ}{dP} = q$ is the Radon-Nikodym derivative. We emphasize that the relative entropy is calculated with respect to reference distribution P. HSTW show that the relative entropy of measure Q with respect to measure P is:

$$\mathcal{R}(Q \parallel P) = \int_0^\infty e^{-\delta t} \mathbb{E}_Q \left[\frac{g_t^2}{2} \right] dt. \tag{5}$$

Given the class of distributions in $\mathcal{Q}$ and a measure of their distance to the true model, HSTW present two classes of preferences which are implicitly defined through control problems. The first class, known as the *constraint* preferences, is specified through the control problem:

$$J^*(\eta) = \sup_{c \in C} \inf_{Q \in \mathcal{Q}(\eta)} \mathbb{E}_Q \left[\int_0^\infty e^{-\delta t} U(c_t, x_t) dt \right] \tag{6}$$

subject to

$$dx_t = \mu(c_t, x_t)dt + \sigma(c_t, x_t)(g_t dt + d\hat{Z}_t) \tag{7}$$

and where

$$\mathcal{Q}(\eta) = \{Q \in \mathcal{Q} : \mathcal{R}(Q \parallel P) \leq \eta\}. \tag{8}$$

The constraint problem appears to be a natural way to study an intertemporal version of Gilboa-Schmeidler. HSTW also introduce a second class of control problems which define the *multiplier* preferences. HSTW are interested in this second class of preferences because they have a natural recursive structure where dynamic programming tools can be employed. Furthermore, HSTW are able to weakly relate the multiplier preferences to the constraint preferences and, hence, weakly relate the multiplier preferences to the underlying theory of choice in Gilboa-Schmeidler. The following control problem implicitly defines the multiplier preferences:

$$\tilde{J}(\theta) = \sup_{c \in C} \inf_{Q \in \mathcal{Q}} \mathbb{E}_Q \left[\int_0^\infty e^{-\delta t} U(c_t, x_t) dt \right] + \theta \mathcal{R}(Q \parallel P) \tag{9}$$

subject to equation (7). Seeing how the preferences are recursive will be useful for our future discussion. A recursive version in discrete time is:

$$V_t = U(c_t) + \beta \inf_{q_t} \left\{ \theta \mathcal{R}_t(Q \parallel P) + \mathbb{E}_t^{q_t}[V_{t+1}] \right\},$$

where V_t is the value function at date t, $\mathcal{R}_t(Q \parallel P)$ is the instantaneous relative entropy, defined as $\mathbb{E}_Q[\log(q_t)]$ and β is the corresponding subjective discount factor. Wang (2001) has axiomatized a class of recursive preferences based on certainty equivalents which look like:

$$V_t = U(c_t) + \beta CE(V_{t+1}),$$

where CE is the certainty equivalent. One valid certainty equivalent is "multi-prior entropy expected utility preferences" presented in Wang (2002). If these static preferences are substituted as the CE, then we obtain the multiplier preferences. Therefore, we can construct an axiomatized version of the multiplier preferences.

In both preference relations, the constants η and θ measure how severely the decision maker penalizes the entropy between the distorted and approximating model and therefore regulates the set of alternative models over which the decision maker wants a robust rule. In HSTW, both η and θ are scalars. The two preference relations are also weakly related: the robustness parameter θ can be thought of as a Lagrange multiplier on the entropy constraint $\mathcal{R}(Q \parallel P) \leq \eta$. However, there is no global connection between the two preferences: HSTW argue that given an η, there exists no θ that makes the two preference orderings agree. The Lagrange multiplier theorem allows for a weak relation between the *local* preferences at the optimum. This relation can be summarized by the following claim established in HSTW: suppose that $\eta = \eta^*, c^*$, and Q^* solve the constraint control problem, then there exists a θ^* such that the multiplier and constraint robust control problems have the same solution. This local result says nothing about preferences orderings off the optimal path.

HSTW give two main reasons for why we should be interested in the multiplier preferences. The first is based on the weak relation that it has to the constraint problem through the Lagrange multiplier theorem. Since the constraint problem appears like Gilboa-Schmeidler, we can utilize the recursive structure of the multiplier preferences to find the optima, and cite the Lagrange multiplier theorem to claim that the optima from the constraint problem would be the same. Therefore the multiplier preferences, through the constraint preferences, are linked to Gilboa-Schmeidler. HSTW give another reason for why the multiplier preferences are interesting: they have been axiomatized by Wang (2002). Wang has also showed that his multi-prior entropy expected utility preferences are consistent with Ellsberg, so if this forms the certainty equivalent, we are analyzing a theory of choice that is consistent with Ellsberg behavior.

The dynamic programming tools given by HSTW to analyze the control problem are a special case of the two-player stochastic differential game analyzed by Fleming and Souganidis (1989). Although the problem is simply a one-person decision problem, it may be fruitful to analyze it as a two-player game. One player of the game is the maximizing agent who chooses an optimal control, while the other player is a minimizing agent choosing the worst model. Thus, HSTW give the following algorithm for computing the optimal policies in the multiplier problem:

$$0 = \max_{c \in C} \min_{g} \Big\{ U(c,x) - \delta V + \frac{\theta}{2} g^2 + [\mu(c,x) + \sigma g] \cdot V_x + \frac{1}{2} \text{trace}[\sigma(c,x)' V_{xx} \sigma(c,x)] \Big\} \tag{10}$$

$$0 = \min_{g} \max_{c \in C} \Big\{ U(c,x) - \delta V + \frac{\theta}{2} g^2 + [\mu(c,x) + \sigma g] \cdot V_x + \frac{1}{2} \text{trace}[\sigma(c,x)' V_{xx} \sigma(c,x)] \Big\}. \tag{11}$$

where HSTW assume the regularity condition needed to reverse the order of operations. At this level of generality, we can solve this program for the worst-case drift distortion to the underlying state. Doing this yields:

$$g_t^* = -\frac{1}{\theta} \sigma(c,x)' V_x,$$

and therefore the distorted state evolves as:

$$dx_t = \left[\mu(c_t, x_t) - \frac{1}{\theta} \sigma(c_t, x_t) \sigma(c_t, x_t)' V_x \right] dt + \sigma(c_t, x_t) d\hat{Z}_t.$$

When the parameter θ approaches $+\infty$, we revert to the ordinary problem because it is too costly for the agent who is minimizing to consider a nonzero perturbation. For small values of θ, there is very little cost to considering a larger set of perturbations. The worst-case distribution also depends on the value function because it is the worst-case calculated at each instant.

We wish to make a few more points about the multiplier preferences. The first is that the multiplier preferences are *not* homothetic. Homotheticity is a very useful property often required for analytical solutions to portfolio choice problems. Recall that V is homothetic if for all $\lambda > 0$, $V(c) \geq V(c') \Leftrightarrow V(\lambda c) \geq V(\lambda c')$ and this condition is clearly not satisfied by the multiplier preferences. Maenhout's transformation is motivated by a desire for homotheticity, so this point is important for the discussion to follow.

Another point has to do with the interpretation of the entropy constraint in HSTW. HSTW are interested in a lifetime entropy constraint, so that the total amount of entropy or discrepancy between the reference model and the worst-case model is less than some number. In the dynamic programming formulation above, the worst-case is calculated at each instant of time, and is a function of the underlying state and value function. Therefore, the HSTW model is considering a different set of perturbations each period and assuming the worst, with the constraint that the cumulative discrepancy is less than some tolerance level.

Since HSTW is largely expressed in terms of control problems, we may gain some intuition for seeing what this means for their interpretation of Gilboa-Schmeidler by thinking of a canonical control problem: a pilot flying a plane. Many applications in the robust control literature are specified as time-zero problems, but thinking of the time t problem of the pilot may be useful. Assume that the pilot's control variable $\{c_t\}$ is the position of his steering wheel and the state of the world $\{x_t\}$ represents the weather and altitude conditions. Let the pilot's goal be to minimize some objective function such as the amount of fuel he uses during the flight.[8] At each instant of time, the pilot does not know what the weather and altitude is outside the cockpit; instead, his sensory equipment is imprecise and only returns a range of values for measuring the state. Suppose that at time t, the plane's equipment registers a reading that it will rain with probability between $50 - 70\%$ and the altitude outside is between $10,000$ and $12,500$ feet. Now, assume that the plane re-

[8] With the additively separable utility specification in HSTW, the agent is interested in minimizing some additively separable function of the fuel used up each instant along the flight.

quires less fuel when there is a lower chance of rain outside and the altitude is higher. Since the pilot's instantaneous objective is to minimize fuel but he is worried that his model is incorrect, the plane's sensor will tell the pilot that the actual environment outside has the probability of rain at 70% and the altitude at $10,000$ feet. At this point, the pilot must make an optimal decision as to where he should move the steering wheel in this worst-case environment. Since the pilot has a lifetime entropy constraint, the next step is to determine how much entropy was just used up at time t. To do so, the plane's entropy monitor will take the worst-case given by the entropy sensor and determine the relative entropy between it and a reference model. The entropy monitor then gives the pilot a reading of the amount of *continuation* entropy he has left to use as he explores the future misspecifications of his weather and altitude monitors. As the pilot travels throughout time, he must keep track of the amount of continuation entropy he has left to use so that the cumulative entropy used up is less than his lifetime constraint.

The point of this exceedingly heuristic exercise is to emphasize some important assumptions. The first is that when the set of distributions is specified by a measure of discrepancy such as the *relative* entropy, it is relative to the reference model. Therefore, while Gilboa-Schmeidler provide a theory of choice for when the decision maker does not know the exact model, delineating the class of distributions with respect to the reference model requires that the decision maker knows the reference model, but still choses to optimize over the worst-case in a set $\mathcal{Q}$. Without knowing the reference model, the pilot above would not be able to calculate the relative entropy. But, if the decision maker *knows* that the reference model is P, why does he *act* as if the model is some worst-case Q? The second important point is that with this lifetime entropy constraint, the worst-case is calculated during each instant of time and can depend on the utility process and various measures of how the utility process changes with the underlying state. In this way, the constraint is like the instantaneous constraint used in the recursive multiple priors model. The third point is that with a lifetime constraint on entropy, the agent must follow the amount of continuation entropy he has left to use. In a sense, he is watching a hidden state variable to ensure that he has used up all of his entropy budget as time progresses. Finally, this framework does not allow the agent to learn about how imprecise the measurement equipment is as time progresses. For instance, it might be natural to suppose that the pilot starts to understand how his devices are imprecise as time progresses and as more observations from these devices become part of his information set.[9]

3. PORTFOLIO CHOICE

The framework of HSTW naturally lends itself to examining Merton's (1969) optimal portfolio choice problem. Maenhout (2001) first used the related framework of Anderson, Hansen, and Sargent (2002) to study the allocation decision between a single risky and riskless asset. We begin by reformulating his discussion in terms of HSTW.[10]

3.1. Maenhout (2001)

An agent seeks to maximize his lifetime utility by optimally selecting the share of his wealth to hold in either a risky asset (α_t) or riskless bond ($1-\alpha_t$). The agent has CRRA preferences of the form:

$$\int_0^T e^{-\delta t}\frac{C_t^{1-\gamma}}{1-\gamma}dt, \tag{12}$$

where $\gamma > 1$ is the coefficient of risk aversion and δ is the discount factor.[11] Fix a probability space $(\Omega, \mathcal{F}, P)$ with filtration $\mathbb{F}$ and let $\{Z_t : t \geq 0\}$ be a one-dimensional adapted Brownian motion. The data generating processes for the riskless asset and the risky asset, respectively, are:

$$\frac{dB_t}{B_t} = rdt, \tag{13}$$

$$\frac{dP_t}{P_t} = \mu dt + \sigma dZ_t. \tag{14}$$

[9]Many other authors have emphasized this drawback of robust control.

[10]The related Anderson, Hansen, and Sargent (2002) framework focuses on misspecification and statistical detection rather than decision theory, while HSTW aims to provide a link to Gilboa and Schmeidler. Adopting this framework is essential to our exercise.

[11]All of our results hold trivially as $\gamma \to 1$. With the appropriate regularity condition, likewise as $T \to \infty$.

The agent's wealth evolves as:

$$dW_t = [W_t(r + \alpha_t(\mu - r)) - C_t]dt + \alpha_t \sigma W_t dZ_t. \tag{15}$$

The two control variables of the agent are α_t and C_t and they satisfy the usual conditions. Merton (1969) showed that the optimal portfolio allocation is $\alpha^* = \frac{1}{\gamma}\frac{\mu - r}{\sigma^2}$, a time and wealth-independent rule.

Adopting the perurbations in HSTW, under measure Q, absolute-continuity implies that a perturbed state equation can be expressed with an added drift term g_t. Such a specification for the set $\mathcal{Q}$ is particularly palatable because Merton (1980) emphasized that an econometrician with access to high-frequency data can estimate a constant variance arbitrarily accurately, so drift distortions will be hard to detect compared with volatility distortions provided that the volatility is constant. Therefore, the perturbed risky asset process is:

$$\frac{dP_t}{P_t} = \mu dt + \sigma(g_t dt + d\hat{Z}_t) \tag{16}$$

and, as a result, the wealth evolves according to:

$$dW_t = [W_t(r + \alpha_t(\mu - r)) - C_t]dt + W_t \alpha_t \sigma(g_t dt + d\hat{Z}_t). \tag{17}$$

Employing the HSTW algorithm to find the optimal policy, we have:

$$0 = \sup_{\alpha_t, C_t} \inf_{g_t} \Big\{ \frac{C_t^{1-\gamma}}{1-\gamma} + [(W_t(r + \alpha_t(\mu - r)) - C_t) + \alpha_t \sigma W_t g_t] \cdot V_w \\ - \delta V + \frac{1}{2}\alpha_t^2 W_t^2 \sigma^2 V_{ww} + \frac{\theta}{2} g_t^2 \Big\}. \tag{18}$$

Solving for the worst-case g_t, we obtain:

$$g_t^* = -\frac{1}{\theta}\alpha_t^* \sigma W_t^* V_w^*. \tag{19}$$

Note this is a transparent adaptation of finding the worst-case program in the abstract presented in the last section. For very large θ, the worst-case model reduces to the benchmark model since nature is significantly penalized for perturbations. The worst-case g_t^* is a function of W_t and the derivative of the value function V_w.

When we substitute g_t^* into (18) and derive the optimality conditions, we find:

$$\partial\alpha_t : \qquad \alpha_t^* = \frac{-V_w}{[V_{ww} - \frac{1}{\theta}V_w^2]W} \cdot \frac{\mu - r}{\sigma^2} \tag{20}$$

$$\partial C_t : \qquad C_t^* = (V_w)^{-\frac{1}{\gamma}}. \tag{21}$$

The only difference from Merton's original optimality condition is the term involving $\frac{1}{\theta}V_w^2$. For large θ, α^* approaches Merton's rule. When we substitute α_t^* and C_t^* to our original problem, we obtain a nonlinear second-order partial differential equation in the value function. This PDE has no well-known solution in general, a fact that is not surprising given that the multiplier preferences are not homothetic.[12] When $\gamma = 1$, there is a closed form solution, but it is well-known that log utility is a very special case in portfolio choice problems.

Maenhout (2001) presents a heuristic argument that the portfolio rule should be homothetic. He specifies that θ be time-dependent and proportional to the value function:

$$\theta_t \equiv \frac{(1-\gamma)V_t}{\lambda}, \tag{22}$$

where λ is a scalar constant and $(1-\gamma)$ is a transformation of this scalar for cosmetic purposes. Recall that in HSTW, θ was a *scalar* which heuristically represented the Lagrangian constraint on the set of alternate models. Under Maenhout's transformation, θ is redefined to be a time-dependent function of the value function.

The central criticism of this paper lies with this transformation. Homotheticity is a crucial element of many economic situations and has played a particularly important role in obtaining analytical solutions in theory of portfolio choice. However, homotheticity is a property of the underlying preferences and therefore should be achieved by modifying this primitive rather than modifying optimality conditions. By changing the interpretation of θ from a scalar to a state-dependent function, we are clearly no longer in the world of Gilboa

[12]If one tries to use the fact that a constant consumption-wealth ratio in the limit requires that the value function take the form $V(W) = \kappa W^{1-\gamma}$ for some constant κ, then there is no closed-form solution.

and Schmeidler. HSTW suggested that θ should be interpreted as a multiplier on the constraint $\mathcal{R}(Q \parallel P) \leq \tau$. The theorem relating the multiplier and constraint preferences at the optima were based on a constant multiplier. In fact, aside from the axiomatization of Wang, the main reason we were interested in the multiplier preferences was because of its connection to the constraint preferences, and by extension Gilboa-Schmeidler. With a new θ_t, we cannot simply cite these results in order to motivate the preferences.

When Maenhout transforms θ_t, the worst case perturbation at the optimum is:

$$g_t^* = -\frac{\lambda}{V_t^*(1-\gamma)}\alpha_t^* \sigma W_t^* V_w^* = -\frac{\lambda}{\gamma+\lambda}\frac{\mu - r}{\sigma}$$

so Maenhout's transformation simply yields a lower drift in the asset price. The worst-case local mean is simply $\mu + \sigma g_t^*$. HSTW insisted on a lifetime entropy constraint so that the set of models changes across time according to endogenous state variables. The point of a lifetime entropy constraint was to ensure that the decision maker was not considering the same set of perturbations each period. But, since the worst case is now independent of time, we have essentially defeated the purpose of a lifetime constraint. With a lower drift, Maenhout's portfolio rule is simply Merton's formula with a lower drift:

$$\alpha_t^* = \frac{1}{\gamma}\frac{[\mu + \sigma g_t^*] - r}{\sigma^2} = \frac{1}{\gamma + \lambda} \cdot \frac{\mu - r}{\sigma^2}.$$

Maenhout, however, interprets robustness as effectively increasing the risk aversion. If we reduced the mean by some constant amount, there would exist a value of λ such that the portfolio rules are identical, so thinking of robustness as an increase in the risk aversion is sensitive to the fact that we are only considering one risky asset.

What happens to the statement of the multiplier problem if we simply substitute the new θ_t to our robust control problem? The problem is not well-defined:

$$\tilde{J}(\lambda) = \sup_{c \in C} \inf_{Q} \mathbb{E}_Q \left[\int_0^\infty e^{-\delta t} U(c_t, x_t) dt \right] + \frac{(1-\gamma)}{\lambda} V_t \cdot \mathcal{R}(Q \parallel P) \quad (23)$$

yet Maenhout still obtains a solution, so the preferences must be defined. We can obtain a better understanding of how they are defined, but not necessarily motivated by Gilboa-Schmeidler, by referring to two other papers utilizing Maenhout's transformation.

3.2. Two other papers

Uppal and Wang (2003) generalize Maenhout (2001) to study portfolio choice over N risky assets and allow for varying degrees of robustness for different components of the multi-dimensional asset price process. The simplest case of preferences are:

$$V_t = u(c_t) + \beta \inf_{\xi} \left[\psi(V_t)\phi L(\xi) + \mathbb{E}_t^{\xi}[V_{t+1}] \right],$$

where V_t is the value function, L is the relative entropy index, ϕ is the weight placed on this index, and ξ is the perturbed density corresponding to q in our discussion above. Uppal and Wang (2003) state that (on page 8): "$\psi(V_t)$ is a *normalization* factor that is introduced to convert the penalty to units of utility so that it is consistent with the units of $\mathbb{E}^{\xi}[V_{t+1}]$; the particular functional form of $\psi(\cdot)$ is often chosen for analytical convenience."[13] Later in the paper, the *normalization* factor is set as $\psi(V_t) = \frac{1-\gamma}{\gamma} V_t$ following Maenhout (2001). The reader is referred to Uppal and Wang (2003) for more details.

In another paper, Liu, Pan, and Wang (2002) study portfolio decisions when an investor is concerned with rare events. They extend Maenhout's model by adding a Poisson jump term to the endowment process. Perturbations then take place over the jump term. Preferences are specified recursively, yet still rely on Maenhout's (2001) homothetic transformation to obtain closed-form solutions. In particular, they express preferences recursively and in discrete-time as:

$$U_t = \frac{c_t^{1-\gamma}}{1-\gamma}\Delta + e^{-\rho\Delta} \inf_{P(\xi)\in\mathcal{P}} \left\{ \frac{1}{\phi}\psi(U_t)\mathbb{E}_t^{\xi}\left[h\left(\ln \frac{\xi_{t+\Delta}}{\xi_t} \right) \right] + \mathbb{E}_t^{\xi}(U_{t+\Delta}) \right\}$$

where ξ is the perturbed density and $h(\cdot)$ is the analog of relative entropy for their environment. They state that (page 7): "following Maenhout (2001), we introduce a *normalization* factor $\psi(U)$ for analytical tractability. To keep the penalty term positive, we let $\psi(U) = (1-\gamma)U$ for $\gamma \neq 1$ and $\psi(U) = 1$ for the log-utility case."[14]

[13] Emphasis added.
[14] Emphasis added.

We apologize for the brief summaries of Uppal and Wang (2003) and Liu, Pan, and Wang (2002), and encourage the reader to refer to both papers. However, since both papers make use of Maenhout's (2001) homothetic transformation which, as we argued in the last section, breaks the link to Gilboa and Schmeidler, these preference specifications are also no longer tied to Gilboa and Schmeidler. We are hesistant to use the terminology adopted by both of these authors that Maenhout's transformation is a *normalization*, because this implies that it is a simple rescaling. Instead, the homothetic transformation dramatically alters the original decision problem. As with Maenhout (2001), it is difficult to specify exactly how problem that both papers solve can be expressed in terms of Gilboa and Schmeidler (1989).

3.3. What problem does Maenhout solve?

Both Uppal and Wang (2003) and Liu, Pan, and Wang (2002) specify preferences recursively and show how preferences may exist even after modified with Maenhout's transformation. Anderson, Hansen, and Sargent (2002), HSTW, Maenhout (2001) and Skiadas (2002) all emphasize the connection between robust control and stochastic differential utility. Anderson, Hansen, and Sargent (2002) showed that the Bellman equation that appears in robust control problem after calculating the worst-case distribution is the same as the Bellman equation with stochastic differential utility. The aggregator (using the terminology of Duffie and Epstein (1992)) is $f(c, V) = U(c) - \delta V$ while the variance multiplier corresponds to θ, a constant in the Anderson, Hansen, and Sargent (2002) framework. In Maenhout (2001), $\theta \propto V_t$, so the variance multiplier is proportional to the value function, a common formulation in stochastic differential utility. Therefore, Maenhout (2001) is able to show that the portfolio rule he derives is observationally equivalent to a Duffie-Epstein investor with elasticity of intertemporal substitution $\frac{1}{\gamma}$ and coefficient of risk aversion $\gamma + \theta$. Skiadas (2002) shows that the equivalence is not only in terms of the Bellman equation, but it is more direct. The multiplier preference relation can be expressed as the solution to the following:

$$dV_t = -\left(U_t - \delta V_t - \frac{1}{2\theta}\sigma_t'\sigma_t\right)dt + \sigma_t' dB_t.$$

When θ is modified as in Maenhout, this preference relation becomes:

$$dV_t = -\left(U_t - \delta V_t - \frac{1}{2\kappa V_t}\sigma_t'\sigma_t\right)dt + \sigma_t' dB_t$$

for some constant κ.

Skiadas (2002) has called such a rescaling a natural generalization. We wish to express a different view in this paper. The rescaling appears natural only when expressed in terms of stochastic differential utility. In fact, Schroder and Skiadas (1999) have already analyzed stochastic differential utility of this form.[15] Our view is that Maenhout's homothetic modification is quite unnatural. One of the aims of the robust control was to construct an intertemporal version of Gilboa-Schmeidler. Showing that the multiplier preferences are isomorphic to stochastic differential utility helps us understand the underlying preferences, but does not bring us any closer to Gilboa-Schmeidler. With Maenhout's transformation, we are four steps from the motivation of Gilboa-Schmeidler. From Gilboa-Schmeidler, we have constructed the constraint preferences, which implied interest in the multiplier preferences, and now preference relations building on Maenhout's transformation have modified the multiplier preferences. Preference relations based on the multiplier control problem were naturally linked to the constraint problem through the observational equivalence result from the Lagrange Multiplier theorem. But, when θ is no longer a constant and is instead a function of the value function, the modified multiplier preferences are no longer linked to the constraint preferences. As a result, they are no longer linked to Gilboa-Schmeidler.

Without any motivating theory of choice or any relation to other preferences, a basic question is why should we be interested in preferences modified by Maenhout's transformation? When expressed as stochastic differential utility the preference relation certainly exists, but what does it have to do the study of robustness or uncertainty

[15]Skiadas (2002) first pointed out the result in Theorem A2 of Schroder and Skiadas (1999).

aversion? If one insists on using Maenhout's transformation, then the onus is to show why Maenhout's transformation is preferred to the multiplier preferences and how the preferences defined by Maenhout's transformation have anything to do with studying robustness or uncertainty aversion. Classes of recursive utility are usually motivated from axiomatized theories of choice such as expected utility, Kreps-Porteus utility, or Chew-Dekel utility. What is the motivating theory of choice for the preferences implicitly defined by Maenhout's transformation? Indeed, to be persuasive, when using Maenhout's transformation, one would need to argue that all existing applications of robust control should be carried out with Maenhout's transformation because it is preferred to the multiplier preferences.

One might wonder if the multiplier should, in fact, be time-varying in some fashion. A new set of preferences defined by a time-varying multiplier may potentially be interesting. However, such a transformation would need to be motivated. With the constant multiplier and constraint binding each period, the goal of HSTW was to allow a different set of perturbations each period in a hands-off manner. Maenhout's time-varying multiplier results in a preference relation where the set of perturbations is the same across time. If our goal is to analyze perturbations of this form, then we can adopt an alternative axiomatized framework which is clearly related to Gilboa-Schmeidler. We describe this alternative in the next section.

4. ALTERNATIVE

Maenhout's (2001) aim in modifying the optimality condition was to obtain a homothetic preference relation. The consequence of his transformation was that the asset price had a lower mean and the set of perturbations the decision maker considered was the same across time. These implications are very similar to the recursive multiple priors framework.[16]

4.1. Recursive multiple-priors

In a special case of the recursive multiple-priors model known as κ-ignorance, we can obtain all of the implications of Maenhout's

[16]See footnote 3 for the relevant citations.

transformation in a portfolio choice context. One of the central differences between the robust control framework and the recursive multiple priors framework is that in recursive multiple priors, the lifetime discrepancy between the reference model and perturbations is not constrained. Instead, recursive multiple priors requires that the set of all measures can be constructed via arbitrary selections from primitive sets of one-step ahead densities. Therefore, recursive multiple priors imposes an instantaneous constraint on the set of conditional one-step ahead densities, while robust control has a lifetime constraint.

There are a number of benefits to thinking of the portfolio choice problem in terms of the recursive multiple priors framework. The first is that it has been axiomatized by Epstein and Schneider (2002b) so the relation to Gilboa-Schmeidler's theory of choice is clear. The second is that the set of priors is clearly delineated. Furthermore, an instantaneous constraint does not require the decision maker to follow the amount of continuation entropy he has left to use up as time progresses. In addition, the framework allows for learning and is not dogmatically pessimistic. Lastly, as a practical matter, the multiplier preferences in HSTW are not homothetic and therefore we should not expect analytical solutions. In contrast, with recursive multiple priors, we can obtain analytical solutions without breaking the connection to Gilboa-Schmeidler.

We can conceptualize the κ-ignorance model of recursive multiple in terms of the robust control formulation in HSTW. Recall the constraint formulation introduced in section 2, where the agent's decision problem is:

$$\inf_{Q \in \mathcal{Q}(\eta)} \mathbb{E}_Q \left[\int_0^\infty e^{-\delta t} U(c_t, x_t) dt \right] \tag{24}$$

such that

$$dx_t = \mu(c_t, x_t)dt + \sigma(c_t, x_t)(g_t dt + d\hat{Z}_t) \tag{25}$$

and

$$\mathcal{Q}(\eta) = \{Q \in \mathcal{Q} : \mathcal{R}(Q \parallel P) \leq \eta\}. \tag{26}$$

Consider this model with a different specification of $\mathcal{Q}(\eta)$. Relative entropy of measure Q with respect to measure P was defined as:

$$\mathcal{R}(Q \parallel P) \equiv \delta \int_0^\infty e^{-\delta t}\mathbb{E}_Q[\log q_t]dt = \int_0^\infty e^{-\delta t}\mathbb{E}_Q\left[\frac{1}{2}g_t^2\right]dt,$$

where q_t is the Radon-Nikodym derivative conditional on $\mathcal{F}_t$. We define the *instantaneous* relative entropy at time t as $\mathcal{R}_t(Q \parallel P) = \mathbb{E}_Q[\log q_t]$, so that:

$$\mathcal{R}(Q \parallel P) = \int_0^\infty e^{-\delta t}\mathcal{R}_t(Q \parallel P)dt. \tag{27}$$

Here, $\mathcal{R}_t(Q \parallel P) = \mathbb{E}_Q\left[\frac{1}{2}g_t^2\right]$. The instantaneous relative entropy is simply the relative entropy of the stochastic process at a particular instant of time.[17]

In the κ-ignorance model, the constraint on the entropy simply restricts the set of models the decision maker explores at each instant of time. Using the language of robust control, the modified constraint problem is:

$$\inf_{Q\in\mathcal{Q}(\tau)} \mathbb{E}_Q\left[\int_0^\infty e^{-\delta t}U(c_t,x_t)dt\right] \tag{28}$$

such that

$$dx_t = \mu(c_t,x_t)dt + \sigma(c_t,x_t)(g_t dt + d\hat{Z}_t) \tag{29}$$

and

$$\mathcal{Q}(\tau) = \{Q \in \mathcal{Q} : \mathcal{R}_t(Q \parallel P) \leq \tau \quad \forall t\}. \tag{30}$$

Since $R_t(Q \parallel P) = \mathbb{E}_Q[\frac{1}{2}g_t^2]$, it is trivial to see that the worst-case perturbation is:

$$g_t^* = -\sqrt{2\tau}. \tag{31}$$

[17]Williams has directed us to two other papers employing the instantaneous constraint in the context of robust control: Lei (2001) and Trojani and Vanini (2002). The aims of our paper and these other papers are quite different, however. In addition to drawing attention to Maenhout's transformation, we also aim to conceptualize recursive multiple priors in terms of robust control. Furthermore, the other two papers are expressed in terms of Anderson, Hansen, and Sargent (2002) and therefore no attempt is made to be explicit about the set of priors or the connection to Gilboa-Schmeidler. Both of these facts are crucial to the discussion here.

For each state of the world, the worst case is implied by the constraint. Furthermore, the worst case g_t^* is independent of time because the constraint is a constant, although nothing prevents us from letting the constant vary with time. The worst case is also independent of the underlying state variables, which is the property that Maenhout (2001) was aiming for with his homothetic transformation. Likewise, the consequence is simply a lower drift on the worst-case asset price:

$$\frac{dP_t}{P_t} = \mu dt + \sigma(g_t dt + d\hat{Z}_t) = (\mu - \sigma\sqrt{2\tau})dt + \sigma d\hat{Z}_t. \tag{32}$$

With these asset price dynamics, the investor simply behaves as a Merton investor facing a lower mean return. Chen and Epstein (2002) first solved the portfolio choice problem for an recursive multiple prior investor. In the next proposition, we present another derivation when the problem is conceptualized as robust control:

Proposition 4.1. *The solution to Merton's optimal portfolio choice problem under the modified constraint formulation given by (28) subject to (29) and (30) is:*

$$\alpha_t^* = \frac{1}{\gamma} \cdot \frac{\mu - r}{\sigma^2} - \frac{1}{\gamma} \cdot \frac{\sqrt{2\tau}}{\sigma} \tag{33}$$

$$C_t^* = \kappa W_t \tag{34}$$

where $\kappa = \kappa(\tau)$ defined in the appendix.

Proof. *See appendix.*[18]

We stress that our contribution in this proposition is to simply highlight the connection to robust control. The portfolio rule can be thought of as Merton's rule with an uncertainty-aversion correction factor: $\frac{1}{\gamma}\frac{\sqrt{2\tau}}{\sigma}$. The stronger constraint formulation implied by this model does not allow the worst-case perturbation to depend on underlying state variables. However, Maenhout's homothetic transformation removes this feature as well. Recall that for the single

[18] All proofs are in the appendix.

asset case Maenhout's portfolio rule was observationally equivalent to an effective increase in risk aversion, while here the portfolio rule is observationally equivalent to an effective decrease in the mean of the risky asset. Clearly, in the one risky asset case, there exists values of τ and θ which imply the same portfolio rule.

4.2. *N* Risky Assets

In this section, we generalize the robust portfolio problem to n assets. Uppal and Wang (2003) first examined this type of generalization for Maenhout's problem using the homothetic transformation. We present another method of generalizing the problem by changing the set of perturbations. The key fact is that the relative entropy of a product measure is the sum of the relative entropy of each component measure. We can get intuition for this fact by recalling that relative entropy is similar to a log-likelihood ratio. Suppose now that the decision maker can choose among n risky assets specified as:

$$\frac{dP_{it}}{P_{it}} = \mu_i dt + \Gamma_i dZ_t, \qquad i = 1..n. \tag{35}$$

where Γ_i is the ith row of an $n \times n$ matrix Γ and $Z = [Z_1, Z_2, ..., Z_n]'$ is a vector of independent Brownian processes. We maintain the restriction that perturbations are absolutely continuous changes of measure. In the n-risky asset case, Girsanov's theorem requires that we simply shift the drift term so that outside a set of Q-measure zero, Z_{it} defined by:

$$Z_{it} = \hat{Z}_{it} + \int_0^t g_{is} ds, \qquad i = 1..n \tag{36}$$

are Brownian motions under Q. The corresponding stock prices follow:

$$\frac{dP_i}{P_i} = \mu_i dt + \Gamma_i (g_t dt + d\hat{Z}_t), \qquad i = 1..n. \tag{37}$$

where $g = [g_1, g_2, ..., g_n]'$. The stochastic process $\{Z_t : t \geq 0\}$ is defined on a space $(\Omega, \mathcal{F})$ with measure $P = P_1 \otimes P_2 \otimes ... \otimes P_n$. Under the perturbed model, a representative probabilty measure is $Q = Q_1 \otimes Q_2 \otimes ... \otimes Q_n \in \mathcal{Q}$.

Proposition 4.2.1: *The relative entropy of the change of measure from P to Q is:*

$$\begin{aligned}\mathcal{R}(Q \parallel P) &= \sum_{i=1}^{n} \mathcal{R}(Q_i \parallel P_i) \\ &= \sum_{i=1}^{n} \int_0^\infty e^{-\delta t}\mathbb{E}_Q\left[\frac{1}{2}g_{it}^2\right]dt.\end{aligned}$$

The proposition above allows us to place n separate penalty terms on each of the perturbations as the decision maker guards against n independent perturbations. This exercise makes the connection to Chen and Epstein (2002) most transparent. Formulated as a robust control problem, the decision problem is:

$$\inf_{Q\in\mathcal{Q}(\tau_1,...,\tau_n)} \mathbb{E}\left[\int_0^\infty e^{-\delta t}U(c_t, x_t)dt\right] \tag{38}$$

subject to

$$dW_t = [W_t(r + \alpha_t(\mu - r)) - C_t]dt + \alpha'\Gamma(g_t dt + d\hat{Z}_t) \tag{39}$$

and

$$\mathcal{Q}(\tau_1, ..., \tau_n) = \{Q \in \mathcal{Q} : \mathcal{R}_t(Q_i \parallel P_i) \leq \tau_i, i = 1..n, \forall t\}. \tag{40}$$

Proposition 4.2.2 *The portfolio rule that optimizes the objective (38) subject to (39) and the constraints is:*

$$\alpha_t^* = \frac{1}{\gamma}(\Gamma\Gamma')^{-1}(\mu - r\iota - \Gamma\Upsilon), \tag{41}$$

$$C_t = \kappa W_t, \tag{42}$$

where $\Upsilon = (\sqrt{2\tau_1}, \sqrt{2\tau_2}, ..., \sqrt{2\tau_n})'$ *and* $\kappa \equiv \kappa(\Upsilon)$ *defined in the appendix.*

Again, Chen and Epstein (2002) first presented the solution for N risky assets and our contribution is simply a robust control derivation. Here, we see that the decision maker acts as if each asset's mean is lower than the approximating model. With the multi-asset generalization, however, the amount that the mean is reduced depends on the matrix Γ and can differ between assets. This weighting differentiates our result from Maenhout (2001) where robustness leads to a higher level of risk aversion.

If $\Gamma_{ij} = 0$ for all $i \neq j$, then the portfolio rule is:

$$\alpha_i = \frac{1}{\gamma} \cdot \frac{\mu_i - r - \Gamma_{ii}\sqrt{2\tau_i}}{\Gamma_{ii}^2} \qquad \text{for } i = 1...n. \tag{43}$$

In the case where the assets are independently distributed, the portfolio rule corresponds to the portfolio rule for each individual asset with a lower mean. While Uppal and Wang (2003) consider different subsets of the risky assets and have penalty terms for each subset, our framework also accomodates this approach.

Uppal and Wang, with Maenhout's adjustment, use the multi-asset generalization to argue that uncertainty aversion may explain the home-bias puzzle and other instances of limited diversification. Their argument harkens back to a (in)famous Keynesian insight:

> *I am in favor of having as large a unit as market conditions will allow... to suppose that safety-first conditions of having a small gamble in a large number of different [companies] where I have no information to reach a good judgement, as compared to a substantial stake in a company where one's information is adequate, strikes me as a travesty of investment policy.*[19]

but has been revisited many times including recent contributions by French and Poterba (1991) and Camerer (1995). It is straightforward to see how the formalized model of uncertainty aversion can generate limited diversification. Consider a two-asset portfolio choice problem, where one asset is the "home" asset and the other asset is the foreign asset. The portfolio rule is:

$$\begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} = \frac{1}{\gamma} \cdot \frac{1}{\sigma_1^2\sigma_2^2 - (\sigma_{12})^2} \cdot \begin{bmatrix} \sigma_2^2\tilde{\mu}_1 - \sigma_{12}\tilde{\mu}_2 \\ \sigma_1^2\tilde{\mu}_2 - \sigma_{12}\tilde{\mu}_1 \end{bmatrix},$$

where $\tilde{\mu}_i = \mu_i - r - \sqrt{2\tau_i}$. Clearly,

$$\frac{\partial \alpha_i}{\partial \tau_i} < 0$$

[19]Quoted from Bernstein (1992).

and for $i \neq j$,

$$\frac{\partial \alpha_i}{\partial \tau_j} > 0, \quad \text{if} \quad \sigma_{12} > 0, \quad \text{and} \quad \frac{\partial \alpha_i}{\partial \tau_j} < 0, \quad \text{if} \quad \sigma_{12} < 0.$$

In words, greater uncertainty aversion with respect to one asset will reduce the corresponding portfolio allocation to that asset. When assets are positively correlated, greater uncertainty aversion with respect to one asset will increase the holdings of the other asset. Epstein (2001) introduces an exchange economy with multi-prior preferences which also rationalizes consumption home bias.

The benefit of this approach is that we are working with a transparent and axiomatized formulation of dynamic Gilboa-Schmeidler. The intuition for both results is similar as uncertainty aversion leads an investor to downweight his portfolio holding in the more uncertain asset. Our aim in showcasing this example is to convince the reader that nothing is lost by using recursive multiple priors over Maenhout's modification of the multiplier preferences.

4.3. Stochastic Differential Utility

Skiadas (2002) emphasizes that stochastic differential utility may lead to a unified perspective on models involving robust control. We find it instructive to write all three preference relations as forms of stochastic differential utility. Chen and Epstein (2002) κ-ignorance is the utility process solving:

$$dV_t = -\Big(U(c_t) - \delta V_t - \kappa \cdot |\sigma_t|\Big) dt + \sigma_t \cdot dZ_t,$$

the HSTW multiplier utility process is defined by:

$$dV_t = -\left(U(c_t) - \delta V_t - \frac{1}{2\theta}\sigma_t' \sigma_t\right) dt + \sigma_t \cdot dZ_t,$$

while the preferences in Maenhout are defined by:

$$dV_t = -\left(U(c_t) - \delta V_t - \frac{1}{2\kappa V_t}\sigma_t' \sigma_t\right) dt + \sigma_t \cdot dZ_t.$$

If we are free to transform the value of the multiplier in the Hansen-Sargent, we can see that it is trivial to also construct the κ-ignorance preferences.

Although all three preference relations appear similar when expressed as stochastic differential utility, relying solely on stochastic differential utility to motivate the preferences obfuscates the point of constructing an intertemporal version of Gilboa and Schmeidler preferences. Indeed, Maenhout's modified preferences only make sense after we have calculated the worst-case and expressed it as stochastic differential utility.[20]

5. OTHER APPLICATIONS

To continue showing that nothing is lost by using recursive multiple priors, we follow Maenhout (2001) to examine the effect of uncertainty aversion when the investor faces stochastic investment opportunities. We first consider a mean-reverting risk premium as in Kim and Omberg (1996). Then we examine the investment opportunity set with jumps in Liu, Pan, and Wang (2002).

5.1. Mean-reverting risk premium

Kim and Omberg (1996) abstract away from a consumption-savings decision to study optimal portfolio choice over terminal wealth. An investor must choose an allocation over a risky asset and a riskfree asset specified, respectively, as:

$$\frac{dP_t}{P_t} = \mu_t dt + \sigma dZ_{1t},$$
$$\frac{dB_t}{B_t} = r dt.$$

Now, we allow the local mean of the risky asset to vary with time. The risk premium is defined as:

$$X_t = \frac{\mu_t - r}{\sigma}$$

which follows a mean-reverting Ornstein-Uhlenbeck process:

$$dX_t = -\lambda(X_t - \bar{X})dt + \sigma_x dZ_{xt}.$$

[20] As Uppal and Wang (2003) point out the connection to Stochastic Differential Utility breaks down when there are varying degrees of uncertainty over the state vector.

We assume that λ, σ_x, and $\bar{X}$ are all positive constants. The Brownian processes are correlated:

$$\mathbb{E}[dZ_{1t}dZ_{xt}] = \rho dt, \quad \rho \in [-1, 1].$$

The investor allocates a fraction α_t of his wealth to the risky asset. Therefore, his wealth follows:

$$dW_t = [W_t(r + \alpha_t(\mu_t - r))]dt + \alpha_t \sigma W_t dZ_{1t}$$

or

$$dW_t = [W_t(r + \alpha_t \sigma X_t)]dt + \alpha_t \sigma W_t dZ_{1t}.$$

A formal statement of the problem of a terminal-wealth investor facing these dynamics is in the appendix. We can establish that:

Proposition 5.1. *The optimal portfolio allocation of the investor is:*

$$\alpha_t^* = \frac{1}{\gamma}\left[\frac{\tilde{X}_t}{\sigma} + (B(t) + C(t)\tilde{X}_t) \cdot \frac{\rho\sigma_x}{\sigma}\right].$$

where the decision-maker believes in a distorted price of risk:

$$\tilde{X}_t = \frac{\mu_t - r - \sigma\sqrt{2\tau_1}}{\sigma}$$

which solves

$$d\tilde{X}_t = -\lambda(\tilde{X}_t - \tilde{\bar{X}})dt + \sigma_x dZ_{xt}$$

with distorted mean

$$\tilde{\bar{X}} = \bar{X} - \frac{\sigma_x}{\lambda}\left(\rho\sqrt{2\tau_1} + \sqrt{1-\rho^2}\sqrt{2\tau_2}\right).$$

$B(t)$ and $C(t)$ form a system of differential equations defined in the appendix.

Notice that the decision maker believes that the price of risk is lower at each instant of time. The decision maker also believes that the long run mean of the risk premium is $\tilde{\bar{X}} < \bar{X}$. The comparative statics of the solutions to the differential equation now simply follow from Kim and Omberg (1996). In particular, when the correlation is negative, the multi-prior investor holds greater amounts of the risky asset, while when the correlation is postive, the investor places less in the risky asset at long horizons than at short horizons. These are the same exact comparative statics as in Maenhout (2001).

5.2. Jumps

The innovation in Liu, Pan, and Wang (2002) (hereafter, LPW) is to add a Poisson jump term to the endowment and hence stock price process. Uncertainty aversion is only specified over the jump process and they analyze the investor's portfolio choice problem and its equilibrium implications. To avoid belaboring the point, we give a concise version of the problem in terms of recursive multiple priors framework to show that Maenhout's transformation is not needed if we employ an instantaenous constraint. Suppose that the endowment follows:

$$dY_t = \mu Y_t dt + \sigma Y_t dZ_t + (e^{Z_t} - 1)Y_{t-} dN_t \tag{44}$$

where N is a Poisson-jump process with intensity λ, and Z_t is a normal random variable with mean μ_J and variance σ_J^2 which controls the jump amplitude. $k \equiv \mathbb{E}[e^Z - 1]$ is the mean percentage jump conditional on a jump.

Consider a stock price which is a deterministic claim on the endowment:

$$P_t = A(t)Y_t,$$

with $A(T) = 0$, so that Ito's Lemma implies:

$$dP_t = \left(\mu + \frac{A'(t)}{A(t)}\right) P_t dt + \sigma P_t dZ_t + (e^{Z_t} - 1)P_{t-} dN_t.$$

We embrace the perturbations examined in LPW, but place an instantaneous constraint on the discrepancy. The alternate measures Q are such that $Q(A) = \mathbb{E}[1_A q_T]$ where the Radon-Nikodym derivative is

$$dq_t = \left(e^{a+bZ_t - b\mu_J - \frac{1}{2}b^2\sigma_J^2} - 1\right) q_{t-} dN_t - (e^a - 1)\lambda q_t dt.$$

LPW show that this perturbation implies that Poisson intensity is $\lambda^q \equiv \lambda e^a$ with mean percentage jump $k^q \equiv (1+k)e^{b\sigma_J^2} - 1$. We saw earlier that the relative entropy was an arbitrary measure of discrepancy. In fact, all that is necessary is a convex function to measure discrepancy. To keep things simple, consider a measure of discrepancy of the form:

$$\mathcal{D}(Q||P) = \begin{bmatrix} (\lambda^q - \lambda)^2 \\ (k^q - k)^2 \end{bmatrix}, \tag{45}$$

so our instantaneous constraint is simply that $\mathcal{D}_1(Q||P) \leq \tau_1$ and $\mathcal{D}_2(Q||P) \leq \tau_2$. The investor must choose consumption and an allocation to the risky stock as before. All other relevant details are in the appendix.

Proposition 5.3 *With an instantaneous constraint, the investor's optimal rule solves:*

$$\left(\mu - r + \frac{1 + A'(t)}{A(t)}\right) - \gamma\alpha^*\sigma^2 + \lambda e^{a^*}\mathbb{E}^{Z(b^*)}\left[(1 + (e^Z - 1)\alpha^*)^{-\gamma}(e^Z - 1)\right] = 0. \tag{46}$$

where (C^*, a^*, b^*) *are defined in the appendix.*

Clearly, we can set equilibrium conditions $C_t = Y_t$ and $\alpha_t = 1$ and derive the equilibrium as in LPW. This exercise will generate the same intuitions as LPW.

6. FURTHER COMMENTS

In this section, we take a step back from both the robust control and recursive multiple priors frameworks to focus on some larger economic questions. We believe that it is an interesting exercise to examine the implications of a non-expected utility theory of choice on portfolio behavior, but we wonder if Gilboa-Schmeidler uncertainty aversion is a convincing alternative. Our discussion will focus on three areas: (1) specifying the set of distributions, (2) using multi-prior decision theory, and (3) empirical relevance.

The applications presented above make clear that the fundamental question with Gilboa and Schmeidler's theory is how to specify the set of alternative distributions $\mathcal{Q}$. This distinction is one of the core differences between robust control and recursive multiple priors. Gilboa and Schmeidler's representation theorem only requires that

$\mathcal{Q}$ is closed and convex, so we suffer from an embarrassment of the riches. By tinkering with this set, we can justify a wide range of behavior. Aware of this potential pitfall, HSTW have emphasized a lifetime constraint on the discrepancy to allow the set of distributions to vary endogenously as the state evolves.

Yet, robust control still allows for a free parameter θ to govern the size of perturbations in the multiplier preferences. We have an analogous situation in Chen and Epstein's (2002) special case, where κ is the free parameter. We saw that moving this parameter around can yield practically any decision rule in a portfolio choice application.[21] Perhaps a boundedly rational model necessitates more parameters, but if we heed Lucas's famous warning to "beware of theorists bearing free parameters" we must ask how we should determine what values of this parameter are reasonable.[22] By "free" parameter, we mean a parameter that is not familiar to us like another parameter of preference. One consistency check is offered by Good (1952). Good argued that min-max rules may be appropriate if the implied worst-case is reasonable. Performing this consistency check will be necessary for every exercise we conduct with these preferences. Perhaps careful experimental methods will also lead to a better understanding of this question. Lastly, one other criteria could be the scope of the model: does it simply explain one fact or does it have more general predictions? Preferences based on Gilboa-Schmeidler have been used to study many of the outstanding puzzles in finance including the equity premium puzzle, limited diversification, limited participation, volatility, and liquidity. Do the free parameters have the same value across these applications? Is it possible to calibrate this parameter without turning the calibration into ex-post rationalization?

In specifying $\mathcal{Q}$, robust control's absolute-continuity formulation requires that the set $\mathcal{Q}$ be "hard to detect" with finite amounts of data. The special case of κ-ignorance in recursive multiple priors implies a similar requirement. With a framework that is based on hard-to-detect perturbations, it must be hard to detect from some reference model. One of our motivations was to examine Knight's

[21]We should also note that this criticism applies to rational Bayesian models of subjective beliefs.

[22]Quoted in Sargent (1993).

uncertainty, when probabilities are unknown, yet we still require the decision maker to know the reference model. Recursive multiple priors does not place as strenuous an assumption on this knowledge as robust control. In recursive multiple priors, the reference model is simply a guide in constructing the set of distributions. In robust control, on the other hand, the decision maker knows the reference model and trusts it enough to calculate the discrepancy relative to it at each step. If a decision maker places so much faith in the reference model, then why is he acting as if his environment is described by the worst-case model? Even though we do not know the number of balls in Ellsberg's ambiguous urn, we must form a reference distribution and calculate entropy with respect to it. This step, in turn, implies that we are priviledging the reference distribution over others or else we would not calculate discrepancy with respect to it. If we have placed this much faith in the reference distribution, why are we unwilling to make it our single prior?

When is it advisable to use a multi-prior decision theory? Some may argue that if it is difficult for the decision maker to specify one subjective prior, it should be no easier for him to specify a set of priors. For Ellsberg's ambiguous urn, however, it may be reasonable to assume a symmetric prior for the ambiguous urn and, thereby, consider a set of distributions. Yet, in this case, a symmetric prior implies indifference and Gilboa-Schmeidler's axioms do not justify Ellsberg's paradox. On the other hand, some may argue that a set of distributions dominates a single distribution because a set gives more flexibility. However, by adopting Gilboa and Schmeidler's theory of choice, our decision is really between a single prior and the worst of a set of priors. A decision maker who is uncomfortable specifying the weights needed to turn a set of priors into a single prior may feel just as uncomfortable relying on the worst distribution among a set.

Another related question is whether a theory of choice based on Gilboa and Schmeidler is normative. This question is particularly relevant for portfolio choice applications where economists usually take on the role of advisors.[23] One may think that uncertainty aversion in Ellsberg's experiment is normative. Likewise, if an agent is

[23]See Campbell and Viceira's (2002) analogy to the Keynesian dentist.

worried about misspecification, it might be logical for him to act as if he had a set of models. But if an agent worried about misspecification adopts the worst-case model, was it really necessary to begin with a set of distributions? Can we study misspecification by simply taking this worst-case and having it serve as our subjective prior? If we believe that this theory is normative, then one implication is that under-diversification in portfolios is normative. Would we recommend that an investor should hold large amounts of company stock in their 401(k) plan? Moreover, a normative intertemporal Gilboa-Schmeidler *must* accomodate learning given that these models imply extreme pessimism. In this regard, recursive multiple priors has made more progress than robust control. Finally, if we eschew a normative notion and adopt only a positive one, then as mentioned earlier, we have a theory that both explains everything and explains nothing.

Sims (2001) has criticized models of robust decision theory in monetary policy settings because he argues that it is not normative for a decision maker to make himself subject to a *dutch book* or *money pump* since the independence axiom is weakened. What happens in the long run to these traders with incorrect beliefs? Kogan, Ross, Wang, and Westefield (2002) have the examined the long run existence and price impact of traders with wrong beliefs, where a wrong belief is simply disagreement over the drift of the stock price process exactly as in the frameworks discussed above. When an agent with wrong beliefs has a portfolio and intermediate consumption decision, they have shown that in the long run these traders do not survive: their relative share of aggregate consumption eventually diminishes. If our Gilboa-Schmeidler agents eventually vanish and have no impact on prices, are they important to study?

Finally, the question of empirical validation is important when examining alternative theories of choice. Limited diversification could be symptomatic of Gilboa-Schmeidler behavior. It could also, however, be symptomatic of a pessimistic subjective prior. While we have formalized the familiarity bias, this formalization does not make it more reasonable as an explanation without a discriminating empirical test. The robust control model allows for a Bayesian interpretation. This feature of their model is distinct from Gilboa-Schmeidler because in Ellsberg's experiment, no single probability measure justifies the paradoxical behavior. If a single pessimistic Bayesian prior can justify behavior, how do we empirically separate it from an agent acting with a set of priors? In the applications we have discussed here, uncertainty simply leads to a lower drift in the asset price. If this is the case, is it possible to separate risk from uncertainty in market environments?

We welcome attempts to build on non-expected utility models, but think that there are number of important questions to think about when using models based on the Gilboa and Schmeidler framework. Similar questions have undoubtedly been asked about other research which studies the implications of weakening or altering the classical expected utility paradigm.

7. CONCLUSION

In this paper we have tried to accomplish three things. First, we studied papers employing Maenhout's transformation and argued that they were poorly motivated from Gilboa-Schmeidler. While the underlying preference relations do exist, we saw that Maenhout's transformation simply makes the worst-case perturbation independent of time and defeats the purpose of the robust control lifetime entropy constraint. In the context of portfolio choice, the transformation simply lead to a lower drift on the asset price. Next, we offered an alternative based on the recursive multiple priors models. In this framework, the set of priors is clearly delineated by an instantaneous constraint and in the special case of κ-ignorance leads to a lower drift on the asset price. We argued that since this framework is axiomatized, transparent, consistent with learning, and yet still leads to analytical solutions, it is the desired framework for studying portfolio choice. We showed how to frame existing applications involving Maenhout's transformation in terms of this model and derive analytical solutions. Finally, we presented some general remarks about this literature and discussed whether Gilboa and Schmeidler is an appropriate non-expected utility model for portfolio choice.

APPENDIX: PROOFS

Proof of 4.1:
At each time t, we have the constraint:

$$\frac{1}{2}\mathbb{E}_Q[g_t^2] \leq \tau. \tag{A.1}$$

To determine the value of g, the value function must satisfy:

$$0 = \sup_{\alpha_t, C_t} \inf_{g_t} \Big\{ \frac{C_t^{1-\gamma}}{1-\gamma} - \delta V + [W_t(r + \alpha_t(\mu - r)) - C_t] \cdot V_w \tag{A.2}$$
$$+ \frac{1}{2}\alpha_t^2 W_t^2 \sigma^2 V_{ww} + \lambda(\frac{1}{2}g_t^2 - \tau) \Big\}$$

with slackness condition:

$$\lambda(\frac{1}{2}g_t^2 - \tau) = 0, \qquad \lambda \geq 0. \tag{A.3}$$

Solving this program for the worst-case g_t, we obtain:

$$g_t^* = -\sqrt{2\tau}. \tag{A.4}$$

Substituting g_t^* into our value function yields:

$$0 = \sup_{\alpha_t, C_t} \Big\{ \frac{C_t^{1-\gamma}}{1-\gamma} - \delta V + [W_t(r + \alpha_t(\mu - r - \sigma\sqrt{2\tau})) - C_t] \cdot V_w \tag{A.5}$$
$$+ \frac{1}{2}\alpha_t^2 W_t^2 \sigma^2 V_{ww} \Big\}$$

with well-known solution:

$$\alpha_t^* = \frac{1}{\gamma} \cdot \frac{\mu - r - \sigma\sqrt{2\tau}}{\sigma^2} \tag{A.6}$$
$$C_t^* = \kappa W_t \tag{A.7}$$

where $\kappa = \kappa(\tau)$ is given by substituting α_t^* and C_t^* into (A.5). After this substitution, we find

$$\kappa(\tau) = \frac{1}{\gamma}\left[\delta - (1-\gamma)r - \frac{(1-\gamma)}{2\gamma}\left(\frac{\mu - r - \sigma\sqrt{2\tau}}{\sigma}\right)^2\right].$$

□

Proof of 4.2.1:
The first part of the proof follows HSTW, which we repeat for clarity. Consider a probability space $(\Omega, \mathcal{F}, P)$ with standard filtration $\mathbb{F}$ and Brownian

motion $\{Z_t : t \geq 0\}$ adapted to this filtration. Let $\mathcal{L}$ be the set of one-dimensional adapted processes. Define the set $\mathcal{L}^2$ by:

$$\mathcal{L}^2 = \left\{ X \in \mathcal{L} : \int_0^T X_t^2 dt < \infty \quad a.s. \right\}. \tag{A.8}$$

A process $g \in \mathcal{L}^2$ satisfies *Novikov's condition* iff:

$$\mathbb{E}\left[\exp\left(\frac{1}{2}\int_0^T g_t^2 dt\right)\right] < \infty \tag{A.9}$$

If the process $g \in \mathcal{L}^2$ satisfies Novikov's condition, then the process q_t defined by:

$$q_t = \exp\left(-\int_0^t g_\tau \cdot dZ_\tau - \frac{1}{2}\int_0^t g_\tau^2 d\tau\right) \tag{A.10}$$

is a martingale. Ito's lemma implies that q_t is an Ito process:

$$dq_t = -q_t g_t dZ_t \tag{A.11}$$

or

$$d(\log q_t) = -g_t \cdot dZ_t - \frac{1}{2} g_t^2 dt. \tag{A.12}$$

We define probability measure Q by:

$$Q(A) = \mathbb{E}[1_A q_T] \tag{A.13}$$

or

$$\frac{dQ}{dP} = q_T. \tag{A.14}$$

Girsanov's theorem states that for a process $g \in \mathcal{L}^2$ such that q is a martingale, the process:

$$\hat{Z}_t = Z_t + \int_0^t g_s ds \tag{A.15}$$

is a Brownian motion under Q. Under Q, we can express the log-likelihood evolution of q_t as:

$$d(\log q_t) = -g_t \cdot d\hat{Z}_t + \frac{1}{2} g_t^2 dt \tag{A.16}$$

In exponential-integral form, we have:

$$q_t = \exp\left[-\int_0^t g_\tau \cdot d\hat{Z}_\tau + \frac{1}{2}\int_0^t g_\tau^2 d\tau\right]. \tag{A.17}$$

The relative entropy of a stochastic process q_t is:

$$\mathcal{R}(Q \parallel P) = \delta \int_0^\infty e^{-\delta t}\mathbb{E}_Q[\log q_t]dt \tag{A.18}$$

$$= \delta \int_0^\infty e^{-\delta t}\mathbb{E}_Q\left[-\int_0^t g_\tau \cdot d\hat{Z}_\tau + \frac{1}{2}\int_0^t g_\tau^2 d\tau\right] dt \tag{A.19}$$

$$= \delta \int_0^\infty \mathbb{E}_Q\left[e^{-\delta t}\int_0^t \frac{g_\tau^2}{2} d\tau\right] dt \tag{A.20}$$

$$= \delta \int_0^\infty \mathbb{E}_Q\left[\frac{g_\tau^2}{2}\right]\int_\tau^\infty e^{-\delta t}dtd\tau \tag{A.21}$$

$$= \int_0^\infty e^{-\delta t}\mathbb{E}_Q\left[\frac{g_t^2}{2}\right] dt. \tag{A.22}$$

Corollary C.3.3 of Dupuis and Ellis (1997) states that if $(\Omega, \mathcal{F})$ is a probability space, and $Q = Q_1 \otimes Q_2$ and $P = P_1 \otimes P_2$ are probability measures on this space, then

$$\mathcal{R}(Q \parallel P) = \mathcal{R}(Q_1 \otimes Q_2 \parallel P_1 \otimes P_2) = \mathcal{R}(Q_1 \parallel P_1) + \mathcal{R}(Q_2 \parallel P_2). \tag{A.23}$$

Since our underlying Brownian motions $\{Z_{1t}, Z_{2t}, ..., Z_{nt} : t \geq 0\}$ are independent, we can perturb each one separately. That is, the Radon-Nikodym theorem states that for each measure Q_i there exists a q_{it} such that:

$$\log q_{it} = -\int_0^t g_{i\tau} \cdot dZ_{i\tau} - \frac{1}{2}\int_0^t g_{i\tau}^2 d\tau \tag{A.24}$$

where Z_{it} for $i = 1...n$ are Brownian motions under P and $\hat{Z}_{it} = Z_{it} + \int_0^t g_{i\tau}d\tau$ for $i = 1...n$ are Brownian motions under Q. Since

$$\mathcal{R}(Q_i \parallel P_i) = \int_0^\infty e^{-\delta t}\mathbb{E}_{Q_i}\left[\frac{g_{it}^2}{2}\right] dt, \qquad \text{for } i = 1...n. \tag{A.25}$$

Therefore,

$$\mathcal{R}(Q \parallel P) = \sum_{i=1}^n \mathcal{R}(Q_i \parallel P_i) = \sum_{i=1}^n \int_0^\infty e^{-\delta t}\mathbb{E}_{Q_i}\left[\frac{g_{it}^2}{2}\right] dt. \tag{A.26}$$

□

Proof of 4.2.2:
Following the argument of Proposition 4.1, we know that:

$$g_{it} = -\sqrt{2\tau_i}, \qquad \text{for } i = 1...n. \tag{A.27}$$

Define $\Upsilon = (\sqrt{2\tau_1}, \sqrt{2\tau_2}, ..., \sqrt{2\tau_n})'$. The perturbed wealth dynamics are:

$$dW_t = [W_t(r + \alpha_t'(\mu - r\iota - \Gamma\Upsilon)) - C_t]dt + \alpha_t'\Gamma dZ_t. \tag{A.28}$$

The value function must satisfy:

$$0 = \sup_{\alpha_t, C_t} \Big\{ \frac{C_t^{1-\gamma}}{1-\gamma} - \delta V + [W_t(r + \alpha_t'(\mu - r\iota - \Gamma\Upsilon)) - C_t] \cdot V_w \tag{A.29}$$
$$+ \frac{1}{2} V_{ww} W_t^2 \alpha_t' \Gamma\Gamma' \alpha_t \Big\}.$$

The first-order conditions are:

$$\partial\alpha: \qquad \alpha_t^* = (\Gamma\Gamma')^{-1}(\mu - r\iota - \Gamma\Upsilon) \cdot \frac{-V_w}{WV_{ww}} \tag{A.30}$$
$$\partial C: \qquad C_t^* = V_w^{-\frac{1}{\gamma}}. \tag{A.31}$$

Setting $V = \kappa^{-\gamma} \frac{W^{1-\gamma}}{1-\gamma}$, we obtain:

$$\alpha_t^* = \frac{1}{\gamma}(\Gamma\Gamma')^{-1}(\mu - r\iota - \Gamma\Upsilon) \tag{A.32}$$
$$C_t^* = \kappa W_t \tag{A.33}$$

where $\kappa = \kappa(\Upsilon)$ is given by substituting α_t^* and C_t^* into (A.29).

□

Proof of 5.1.
To perturb the assets and obtain a measure of the perturbation, we find it convenient to work with independent Brownian motions. Define $\{Z_{2t} : t \geq 0\}$ to be a Brownian motion independent of $\{Z_{1t} : t \geq 0\}$ so that

$$dZ_{xt} = \rho dZ_{1t} + \sqrt{1-\rho^2} dZ_{2t}. \tag{A.34}$$

As before, we can perturb each of these Brownian motions separately to obtain:

$$dZ_{1t} = g_{1t}dt + d\hat{Z}_{1t}, \tag{A.35}$$
$$dZ_{2t} = g_{2t}dt + d\hat{Z}_{2t}, \tag{A.36}$$

and therefore

$$dZ_{xt} = \left[\rho g_{1t} + g_{2t}\sqrt{1-\rho^2}\right] dt + \rho d\hat{Z}_{1t} + \sqrt{1-\rho^2} d\hat{Z}_{2t}. \tag{A.37}$$

The investor's decision problem is:

$$\sup_{\alpha} \inf_{Q \in \mathcal{Q}(\tau_1, \tau_2)} \mathbb{E}_Q \left[\int_0^\infty e^{-\delta t} U(W_t) dt \right]$$

such that

$$dW_t = [W_t(r + \alpha_t \sigma X_t)]dt + \sigma W_t (g_{1t} dt + d\hat{Z}_{1t}),$$

$$dX_t = -\lambda(X_t - \bar{X})dt + \sigma_x[(\rho g_{1t} + g_{2t}\sqrt{1-\rho^2})dt + \rho d\hat{Z}_{1t} + \sqrt{1-\rho^2} d\hat{Z}_{2t}],$$

and

$$\mathcal{Q}(\tau_1, \tau_2) = \{Q \in \mathcal{Q} : \mathcal{R}_t(Q_1 \parallel P_1) \leq \tau_1, \mathcal{R}_t(Q_2 \parallel P_2) \leq \tau_2 \quad \forall t\}$$

As we know from previous example, the worst case will simply be $g_{1t} = -\sqrt{2\tau_1}$ and $g_{2t} = -\sqrt{2\tau_2}$. The most intuitive way to solve this problem is to define a distorted price of risk:

$$\tilde{X}_t \equiv \frac{\mu_t - r - \sigma\sqrt{2\tau_1}}{\sigma} = X_t - \sqrt{2\tau_1}, \tag{A.38}$$

with distorted long-term mean:

$$\tilde{\bar{X}} = \bar{X} - \frac{\sigma_x}{\lambda}\left(\rho\sqrt{2\tau_1} + \sqrt{1-\rho^2}\sqrt{2\tau_2}\right). \tag{A.39}$$

Therefore, we can phrase the state-evolution in terms of the distorted risk premium:

$$d\tilde{X}_t = -\lambda(\tilde{X}_t - \tilde{\bar{X}})dt + \sigma_x d\hat{Z}_{xt}$$

and

$$dW_t = [W_t(r + \alpha_t \sigma \tilde{X}_t)]dt + \alpha_t \sigma W_t d\hat{Z}_{1t}$$

where $d\hat{Z}_{xt} = \rho d\hat{Z}_{1t} + \sqrt{1-\rho^2} d\hat{Z}_{2t}$.

With this adjustment, we can simply recall Kim and Omberg's (1996) solution:

$$\alpha_t^* = \frac{1}{\gamma}\left[\frac{\tilde{X}_t}{\sigma} + (B(t) + C(t)\tilde{X}_t) \cdot \frac{\rho\sigma_x}{\sigma}\right]. \tag{A.40}$$

where

$$\frac{dC}{dt} = cC^2(t) + bC(t) + a, \tag{A.41}$$

$$\frac{dB}{dt} = cB(t)C(t) + \frac{b}{2}B(t) + \lambda_x \tilde{\bar{X}} C(t), \tag{A.42}$$

$$\frac{dA}{dt} = \frac{c}{2}B^2(t) + \frac{1}{2}\sigma_x^2 C(t) + \lambda_x \tilde{\bar{X}} B(t), \tag{A.43}$$

$$C(0) = B(0) = C(0) \tag{A.44}$$

with $a = \frac{1}{\gamma} - 1$, $b = 2[(\frac{1}{\gamma} - 1)\rho\sigma_x - \lambda_x]$, and $c = \sigma_x^2[1 + \frac{1}{\gamma}]$.

□

Proof of 5.2.
Since $\mathcal{D}_1(Q||P) \leq \tau_1$ and $\mathcal{D}_2(Q||P) \leq \tau_2$, in the worst-case, both constraints bind, yielding:

$$a^* = \log\left(\frac{1}{\lambda}\sqrt{\tau_1} + 1\right) \tag{A.45}$$

and

$$b^* = \frac{1}{\sigma_J^2}\log\left(\frac{\tau_2}{(1+k)^2} + 1\right) \tag{A.46}$$

This implies the worst-case jump intensity increases to $\lambda + \lambda\sqrt{\tau_1}$ while the mean increases to $(1+k) + \frac{\tau_2}{1+k}$. As in all previous examples, we have two free parameters (τ_1, τ_2) which can lead to any pessimistic twisting that we want.

The investor's wealth solves:

$$dW_t = \left\{\left[r + \alpha_t\left(\mu - r + \frac{1 + A'(t)}{A(t)}\right)\right] W_t - C_t\right\} dt + \alpha_t \sigma W_t dZ_t$$
$$+\alpha_{t-} W_{t-}(e^{Z_t} - 1) dN_t \tag{A.47}$$

With the worst-case a^* and b^*, the Bellman equation is:

$$0 = \sup_{\alpha_t, C_t}\left\{u(c) - \rho J + J_t + \left[r + \alpha_t\left(\mu - r + \frac{1 + A'(t)}{A(t)}\right)\right]\right.$$
$$+WJ_w - CJ_w + \frac{\sigma^2}{2}\alpha^2 W^2 J_{ww}$$
$$\left. +\lambda e^{a^*}\left(\mathbb{E}^{Z(b^*)}[J(W(1 + \alpha(e^Z - 1)))] - J(W,t)\right)\right\} \tag{A.48}$$

where $\mathbb{E}^{Z(b)}[f(Z)] = \mathbb{E}[e^{(bZ - b\mu_J - \frac{1}{2}b^2\sigma_J^2)} f(Z)]$. Suppose that the value function is multiplicatively separable:

$$J(W,t) = \frac{W^{1-\gamma}}{1-\gamma} f(t)^\gamma$$

The optimal consumption is then $C_t^* = f(t)^{-1} W_t$ and the optimal portfolio weight solves:

$$\left(\mu - r + \frac{1 + A'(t)}{A(t)}\right) - \gamma\alpha^*\sigma^2 + \lambda e^{a^*}\mathbb{E}^{Z(b^*)}\left[(1 + (e^Z - 1)\alpha^*)^{-\gamma}(e^Z - 1)\right] = 0. \tag{A.49}$$

Finally, $f(t)$ is the solution to:

$$\frac{\gamma}{1-\gamma}\frac{1+f'(t)}{f(t)} - \frac{\rho}{1-\gamma} + r + \alpha_t \left(\mu - r + \frac{1+A'(t)}{A(t)}\right) - \frac{1}{2}\gamma\sigma^2\alpha^2$$
$$+\frac{1}{1-\gamma}\lambda e^{a^*}\left(\mathbb{E}^{Z(b^*)}[(1+(e^Z-1)\alpha^*)^{1-\gamma}] - 1\right) = 0. \quad \text{(A.50)}$$

Therefore we have solved for $(C^*, \alpha^*, a^*, b^*)$.

□

REFERENCES

Anderson, E., L. Hansen, and T. Sargent (2002): "A Quartet of Semi-Groups for Model Specification, Detection, Robustness, and the Price of Risk," Working paper, University of Chicago, UNC, and Stanford University.

Anscombe, F., and R. Aumann (1963): "Definition of Subjective Probability," *Annals of Mathematical Statistics*, 34, 199–205.

Bernstein, P. L. (1992): *Capital Ideas: The Improbable Origins of Modern Wall Street*. New York: Free Press.

Camerer, C. (1995): "Individual Decision Making," in *Handbook of Experimental Economics*, ed. by A. Roth, and J. Kagel, pp. 587–703. Princeton University Press.

Campbell, J. Y., and L. Viceira (2002): *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press.

Casadesus-Masanell, R., P. Klibanoff, and E. Ozdenoren (2000): "Maxmin Expected Utility over Savage Acts with a Set of Priors," *Journal of Economic Theory*, pp. 35–65.

Chamberlain, G. (2000): "Econometric Applications of Maxmin Expected Utility," *Journal of Applied Econometrics*.

Chen, Z., and L. Epstein (2002): "Ambiguity, Risk and Asset Returns in Continuous Time," *Econometrica*, 70, 1403–1443.

Duffie, D., and L. Epstein (1992): "Stochastic Differential Utility," *Econometrica*, 60, 353–394.

Dupuis, P., and R. Ellis (1997): *A Weak Convergence Approach to the Theory of Large Deviations*. Wiley Series in Probability and Statistics. New York: John Wiley and Sons.

Ellsberg, D. (1961): "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics*, 75, 643–669.

Epstein, L. (2001): "Sharing Ambiguity," *American Economic Review*, 91, 45–50.

Epstein, L., and M. Schneider (2002a): "Learning Under Ambiguity," Unpublished paper, University of Rochester and UCLA.

——— (2002b): "Recursive Multiple Priors," Unpublished paper, University of Rochester and UCLA.

Epstein, L., and T. Wang (1994): "Intertemporal Asset Pricing Under Knightian Uncertainty," *Econometrica*, 62, 283–322.

Epstein, L., and S. Zin (1989): "Substitution, Risk Aversion and the Temporal Behaviro of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–969.

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.

French, K., and J. Poterba (1991): "Investor Diversification and International Equity Markets," *American Economic Review*, 81, 222–226.

Gilboa, I., and D. Schmeidler (1989): "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, 18, 141–53.

Good, I. J. (1952): "Rational Decisions," *Journal of the Royal Statistical Society: Series B*, 14, 107–114.

Hansen, L., and T. Sargent (2001a): "Robust Control and Filtering for Macroeconomics," Unpublished manuscript, University of Chicago and Stanford University.

——— (2001b): "Robust Control and Model Uncertainty," *American Economic Review*, 91, 60–66.

——— (2001c): "Time Inconsistency of Robust Control?," Unpublished paper, University of Chicago and Stanford University.

——— (2001d): "Wanting Robustness in Macroeconomics," Unpublished paper, University of Chicago and Stanford University.

Hansen, L., T. Sargent, G. Turmuhambetova, and N. Williams (2002): "Robust Control and Model Uncertainty," Unpublished paper, University of Chicago and Stanford University.

Kim, T. S., and E. Omberg (1996): "Dynamic Nonmyopic Portfolio Behavior," *Review of Financial Studies*, 9, 141–161.

Klibanoff, P. (1995): "Dynamic Choice and Uncertainty Aversion," Unpublished paper, Northwestern MEDS.

Knight, F. (1921): *Risk, Uncertainty, and Profit*. Houghton Mifflin Company.

Knox, T. (2002): "Learning How to Invest when Returns are Uncertain," Unpublished paper, Harvard University.

Kogan, L., S. Ross, J. Wang, and M. Westefield (2002): "The Survival and Price Impact of Irrational Traders," Unpublished paper, MIT.

Lei, C. (2001): "Why Don't Investors Have Large Positions in Stocks? A Robustness Perspective," Unpublished paper, University of Chicago.

LIU, J., J. PAN, AND T. WANG (2002): "An Equilibrium Model of Rare Event Premia," Unpublished paper, MIT.

MAENHOUT, P. (2001): "Robust Portfolio Rules, Hedging and Asset Pricing," Unpublished paper, INSEAD.

MERTON, R. C. (1969): "Lifetime Portfolio Selection Under Uncertainty," *Review of Economics and Statistics*, 51, 247–257.

——— (1980): "On Estimating the Expected Return on the Market: An Exploratory Investigation," *Journal of Financial Economics*, 8, 323–361.

SARGENT, T. (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press.

SCHRODER, M., AND C. SKIADAS (1999): "Optimal Consumption and Portfolio Choice with Stochastic Differential Utility," *Journal of Economic Theory*, 89, 68–126.

SIMS, C. (2001): "Pitfalls of a Minmax Approach to Model Uncertainty," *American Economic Review*, 91, 51–54.

SINISCALCHI, M. (2001): "Bayesian Updating for General Maxmin Expected Utility Preferences," Unpublished paper, Princeton University.

SKIADAS, C. (2002): "Robust Control and Recursive Utility," Unpublished paper, Northwestern University.

TROJANI, F., AND P. VANINI (2002): "A Note on Robustness in Merton's Model of Intertemporal Consumption and Portfolio Choice," *Journal of Economic Dynamics and Control*, 26, 423–435.

UPPAL, R., AND T. WANG (2003): "Model Uncertainty and Underdiversification," *Journal of Finance*, forthcoming.

WANG, T. (2001): "Conditional Preferences and Updating," forthcoming, Journal of Economic Theory.

——— (2002): "Two Classes of Multi-Prior Preferences," Unpublished paper, University of British Columbia.