Robust Decision Theory and Econometrics

Gary Chamberlain[1]

[1]Department of Economics, Harvard University, Cambridge, MA 02138, USA; gary_chamberlain@harvard.edu (13 January 2020)

Xxxx. Xxx. Xxx. Xxx. YYYY. AA:1–33

https://doi.org/10.1146/((please add article doi))

Keywords

portfolio choice, subjective expected utility, maxmin expected utility, variational preferences, constraint preferences, smooth ambiguity preferences

Abstract

The paper uses the empirical analysis of portfolio choice to illustrate econometric issues that arise in decision problems. Subjective expected utility (SEU) can provide normative guidance to an investor making a portfolio choice. The investor, however, may have doubts on the specification of the distribution and may seek a decision theory that is less sensitive to the specification. I consider three such theories: maxmin expected utility, variational preferences (including multiplier and divergence preferences and the associated constraint preferences), and smooth ambiguity preferences. A simple two-period model is used to illustrate their application. Normative empirical work on portfolio choice is mainly in the SEU framework and bringing in ideas from robust decision theory may be fruitful.

1. INTRODUCTION

I shall use the empirical analysis of portfolio choice to illustrate econometric issues that arise in decision problems. For example, Barberis (2000) considers an investor with power utility over terminal wealth. There are two assets: Treasury bills and a stock index. The investor uses a vector autoregression to specify a likelihood function, and he specifies a prior distribution on the parameters. Barberis considers both a buy-and-hold strategy and dynamic rebalancing. The empirical work uses postwar data on asset returns and the dividend yield. The framework is Bayesian decision theory, corresponding to subjective expected utility (SEU) preferences, as in Ramsey (1926), Savage (1954), and Anscombe & Aumann (1963). I believe this is the correct normative framework for decision making under uncertainty. Nevertheless, the investor may have doubts about his model; that is, about the predictive distribution for future returns (which may have been based on specifying a likelihood function and a prior). Diaconis & Skyrms (2018, p. 43) note that Ramsey was aware that the assumptions going into utility-probability representations are highly idealized, and they provide the following quote from his 1926 paper on "Truth and Probability":

> I have not worked out the mathematical logic of this in detail, because this would, I think, be rather like working out to seven places of decimals a result only valid to two. My logic cannot be regarded as giving more than the sort of way it might work.

It is difficult to specify a predictive distribution precisely. Hansen & Sargent (2001, 2008, 2015) respond to this difficulty by using robust control theory, which they relate to work on ambiguity in decision theory, including Gilboa and Schmeidler (1989), Maccheroni et al. (2006a), Klibanoff et al. (2005), and Strzalecki (2011). Hansen & Sargent achieve robustness by working with a neighborhood of the reference model and maximizing the minimum of expected utility over that neighborhood.

The Ellsberg (1961) paradox has played a key role in developing alternatives to SEU in decision theory. The following example is from Klibanoff et al. (2005). Table 1 shows four acts: f, g, f', and g', with payoffs contingent on three (mutually exclusive and exhaustive) events: A, B, and C. This could correspond to an urn with 30 balls of color A and 60 balls divided in some unknown way between colors B and C. The decision maker is asked to rank f and g, and to rank f' and g'. Savage's axiom P2, the sure-thing principle, states that if two acts are equal on a given event, then it should not matter (for ranking the acts in terms of preferences) what they are equal to on that event. The payoffs to f and g are 0 if C occurs. The ranking should not change if instead that payoff is 10. So if f is preferred to g, then f' should be preferred to g'. I believe this is the correct normative conclusion. Nevertheless, one can argue for $f \succ g$ because $E[u(f)] = \frac{1}{3}u(10) + \frac{2}{3}u(0)$, whereas evaluating $E[u(g)]$ requires assigning a probability to the event B, when we are told only that it is between 0 and 2/3. Likewise, one can argue for $g' \succ f'$ because $E[u(g')] = \frac{1}{3}u(0) + \frac{2}{3}u(10)$, whereas evaluating $E[u(f')]$ requires assigning a probability to the event C, when we are told only that it is between 0 and 2/3.

One motivation for alternatives to SEU preferences is the positive one of modeling observed behavior, where $f \succ g$ and $g' \succ f'$ are common choices. My interest in alternatives to SEU is, however, normative, in seeking a more robust decision theory that is less sensitive to model specification.

Section 2 uses a simple two-period portfolio choice problem to present alternatives to SEU. The applications of robust decision theory in Hansen & Sargent are mainly in general

Table 1 Ellsberg Example

	A	B	C
f	10	0	0
g	0	10	0
f'	10	0	10
g'	0	10	10

equilibrium problems in macro economics and asset pricing. An alternative application is to portfolio choice. Section 3 discusses empirical work on normative portfolio choice, in which the investor takes the distribution of asset prices as given. This work is mainly in the SEU framework, and I think that bringing in ideas from the ambiguity literature may be fruitful. Section 4 considers how these ideas could be applied to questions in Barberis (2000). Section 5 considers a normative critique of ambiguity preferences. Section 6 concludes.

2. PORTFOLIO CHOICE: THEORY

This section develops basic concepts in a simple setting. We begin with a single prior Bayesian decision analysis based on subjective expected utility (SEU) preferences. Next we consider the multiple priors framework of Gilboa & Schmeidler (1989), and the special case of a rectangular set of priors developed by Epstein & Schneider (2003). Then we examine multiplier preferences, which are a special case of the variational preferences in Maccheroni et al. (2006a). Strzalecki (2011) provides axioms that characterize multiplier preferences. In some cases, multiplier preferences have a recursive form. Divergence preferences are an important class of variational preferences, and they include multiplier preferences as a special case. We connect divergence preferences to constraint preferences, which are a special case of the multiple priors framework. Then we apply the smooth ambiguity preferences of Klibanoff et al. (2005).

2.1. Single Prior

Consider a two period problem with $t = 0, 1$. Investment decisions are made at the beginning of each period. Initial wealth is given: $W_0 = w_0 > 0$. Utility is of the power form over (random) final wealth:

$$u(W_2) = \begin{cases} \frac{W_2^{1-\gamma}}{1-\gamma} & \text{if } \gamma \neq 1, \\ \log(\gamma) & \text{if } \gamma = 1. \end{cases}$$

There is one riskless asset with a (gross) return r_f. There is one risky asset; its gross return in period t is R_t, which can take on two values h (high) and l (low). Assume that $0 < l < r_f < h$. The investor treats R_0 and R_1 as exchangeable and specifies the following likelihood function and prior distribution: conditional on θ, R_0 and R_1 are independent and identically distributed with

$$\Pr(R_t = h \,|\, \theta) = \theta, \quad \Pr(R_t = l \,|\, \theta) = 1 - \theta \quad (t = 0, 1).$$

The prior for θ is a beta distribution with parameters α and β; the density function is

$$\frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}\theta^{\alpha-1}(1-\theta)^{\beta-1}.$$

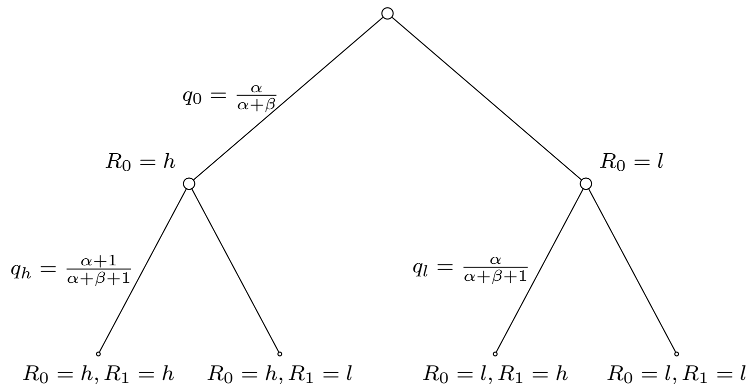

Figure 1

Probability tree for the predictive distribution in (1) implied by (R_0, R_1) i.i.d. conditional on θ, with a beta distribution for θ with parameters α and β.

The implied predictive distribution for (R_0, R_1) can be factored as follows:

$$\begin{aligned} \Pr(R_0 = h) &= E(\theta) = \frac{\alpha}{\alpha+\beta} = q_0, \qquad (1)\\ \Pr(R_1 = h \mid R_0 = h) &= E(\theta \mid R_0 = h) = \frac{\alpha+1}{\alpha+\beta+1} = q_h,\\ \Pr(R_1 = h \mid R_0 = l) &= E(\theta \mid R_0 = l) = \frac{\alpha}{\alpha+\beta+1} = q_l. \end{aligned}$$

Let q denote this predictive distribution. See Figure 1.

Let a_0 denote the fraction of wealth that is invested in the risky asset in period 0. In period 1, the investor has observed R_0; the portfolio weight on the risky asset equals $a_1(h)$ if $R_0 = h$ and equals $a_1(l)$ if $R_0 = l$. Simplify notation by setting $a' = (a_0, a_1(l), a_1(h))$ and denote $0 \le a_0, a_1(l), a_1(h) \le 1$ by $0 \le a \le 1$. The investor has the following problem:

$$\max_{0\le a\le 1} E_q\Big[u\Big(w_0[(R_0 - r_f)a_0 + r_f][(R_1 - r_f)a_1(R_0) + r_f]\Big)\Big]. \qquad (2)$$

The investor's preferences are recursive, and this problem can be solved by dynamic programming (as in Kreps & Porteus (1979)). Using iterated expectations, the objective function is

$$\begin{aligned} &E_q\Big[E_q\Big[u\Big(w_0[(R_0 - r_f)a_0 + r_f][(R_1 - r_f)a_1(R_0) + r_f]\Big) \mid R_0\Big]\Big]\\ &\quad \le E_q\Big[\max_{0\le x\le 1} E_q\Big[u\Big(w_0[(R_0 - r_f)a_0 + r_f][(R_1 - r_f)x + r_f]\Big) \mid R_0\Big]\Big]. \end{aligned}$$

Define the maximized value functions

$$\begin{aligned} J_h(w) &= \max_{0\le x\le 1}\Big[q_h u\Big(w[(h - r_f)x + r_f]\Big) + (1 - q_h)u\Big(w[(l - r_f)x + r_f]\Big)\Big],\\ J_l(w) &= \max_{0\le x\le 1}\Big[q_l u\Big(w[(h - r_f)x + r_f]\Big) + (1 - q_l)u\Big(w[(l - r_f)x + r_f]\Big)\Big]. \end{aligned}$$

Then the optimal portfolio weight on the risky asset in period 0 is

$$\begin{aligned} a_0^* &= \arg\max_{0\le x\le 1} E_q\Big[J_{R_0}\Big(w_0[(R_0 - r_f)x + r_f]\Big)\Big] \qquad (3)\\ &= \arg\max_{0\le x\le 1}\Big[q_0 J_h\Big(w_0[(h - r_f)x + r_f]\Big) + (1 - q_0)J_l\Big(w_0[(l - r_f)x + r_f]\Big)\Big]. \end{aligned}$$

2.2. Multiple Priors

Gilboa & Schmeidler (1989) use the Anscombe & Aumann (1963) framework, which distinguishes between a roulette lottery, in which probabilities are given, and a horse lottery, in which probabilities are not given. In the Ellsberg problem in Table 1, a bet on A corresponds to a roulette lottery, with $\Pr(A) = 1/3$, whereas a bet on B corresponds to a horse lottery, where we are given only that $0 \leq \Pr(B) \leq 2/3$. Let the set Z denote the possible consequences (outcomes, prizes), and let $\triangle(Z)$ denote probability distributions on Z with finite support (roulette lotteries). An act (horse lottery, payoff profile) f is a finite-valued mapping from the state space S to lotteries over consequences: $f\colon S \to \triangle(Z)$; the set of all such acts is denoted $\mathcal{F}(\triangle(Z))$. Gilboa & Schmeidler weaken the independence axiom in Anscombe & Aumann to certainty independence: for all $f, g \in \mathcal{F}(\triangle(Z))$ and $h \in \triangle(Z)$, and for all $\alpha \in (0,1)$, $f \succ g$ if and only if $\alpha f + (1-\alpha)h \succ \alpha g + (1-\alpha)h$. (The convex combination $\alpha f + (1-\alpha)h$ is an act whose value at $s \in S$ is the probability mixture of roulette lotteries $\alpha f(s) + (1-\alpha)h$.) The Anscombe & Aumann independence axiom, which leads to SEU, holds for all acts $h \in \mathcal{F}(\triangle(Z))$, not just for roulette lotteries $h \in \triangle(Z)$.

With Gilboa & Schmeidler preferences, there is a closed, convex set $\mathcal{C}$ of prior distributions on S. An investor with these preferences evaluates a portfolio plan by calculating expected utility under each prior $q \in \mathcal{C}$ and then taking the minimum. So the investor's problem is

$$\max_{0\leq a\leq 1} \min_{q\in\mathcal{C}} E_q\Big[u\Big(w_0[(R_0 - r_f)a_0 + r_f][(R_1 - r_f)a_1(R_0) + r_f]\Big)\Big]; \tag{4}$$

that is, maxmin expected utility (MEU). For example, let q_k be the predictive distribution for (R_0, R_1) in (1), with parameters α_k and β_k $(k = 1, \ldots, K)$. Then we can let $\mathcal{C}$ consist of all probability mixtures:

$$\mathcal{C} = \{\sum_{k=1}^{K} \xi_k q_k : \xi_k \geq 0, \sum_{k=1}^{K} \xi_k = 1\}. \tag{5}$$

The minimax theorem can be used to reverse the order of minimization and maximization in (4). Then, for a given q, the inner maximization is the investor's problem with a single prior as in (2). Let $V(q)$ denote the maximized value. Then the least-favorable prior is

$$q^* = \arg\min_{q\in\mathcal{C}} V(q).$$

Chamberlain (2000b) develops an algorithm for solving this problem, based on a convex program. Once the least favorable prior q^* has been obtained, the investor solves the single prior problem in (2) using q^*.

2.2.1. Rectangular Set of Priors. I shall follow Knox (2003a, b) in using the two-period model to examine the restrictions implied by the Epstein & Schneider (2003) condition that the convex set $\mathcal{C}$ be rectangular. This condition takes the following form:

$$\begin{aligned} \Pr(R_0 = h) &\in [\underline{q}^0, \overline{q}^0], \\ \Pr(R_1 = h \,|\, R_0 = h) &\in [\underline{q}^h, \overline{q}^h], \\ \Pr(R_1 = h \,|\, R_0 = l) &\in [\underline{q}^l, \overline{q}^l]. \end{aligned}$$

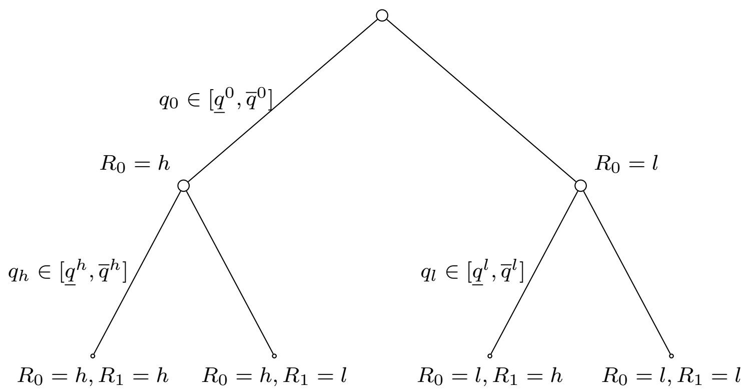

Figure 2

Probability tree for the rectangular set of prior distributions for (R_0, R_1).

The key here is that in forming the set of distributions $\mathcal{C}$ for (R_0, R_1), any marginal distribution for R_0 can be combined with any conditional distribution for R_1 given R_0. So select any three values: $q_0 \in [\underline{q}^0, \overline{q}^0]$, $q_h \in [\underline{q}^h, \overline{q}^h]$, and $q_l \in [\underline{q}^l, \overline{q}^l]$. Then the following distribution is an element of $\mathcal{C}$:

$$\begin{aligned}
\Pr(R_0 = h, R_1 = h) &= q_0 q_h, \\
\Pr(R_0 = h, R_1 = l) &= q_0(1 - q_h), \\
\Pr(R_0 = l, R_1 = h) &= (1 - q_0) q_l, \\
\Pr(R_0 = l, R_1 = l) &= (1 - q_0)(1 - q_l).
\end{aligned}$$

See Figure 2.

When $\mathcal{C}$ is rectangular, we can use iterated expectations to break up the minimization over $\mathcal{C}$ into three separate minimizations:

$$\begin{aligned}
\min_{q_0 \in [\underline{q}^0, \overline{q}^0]} \Big[q_0 \min_{q_h \in [\underline{q}^h, \overline{q}^h]} \Big[& q_h u\Big(w_0[(h - r_f)a_0 + r_f][(h - r_f)a_1(h) + r_f] \Big) \\
& + (1 - q_h) u\Big(w_0[(h - r_f)a_0 + r_f][(l - r_f)a_1(h) + r_f] \Big) \Big] \\
+ (1 - q_0) \min_{q_l \in [\underline{q}^l, \overline{q}^l]} \Big[& q_l u\Big(w_0[(l - r_f)a_0 + r_f][(h - r_f)a_1(l) + r_f] \Big) \\
& + (1 - q_l) u\Big(w_0[(l - r_f)a_0 + r_f][(l - r_f)a_1(l) + r_f] \Big) \Big] \Big].
\end{aligned}$$

The preferences are recursive and the investor's problem can be solved by dynamic programming:

$$\begin{aligned}
J_h(w) &= \max_{0 \le x \le 1} \min_{q_h \in [\underline{q}^h, \overline{q}^h]} \Big[q_h u\Big(w[(h - r_f)x + r_f] \Big) + (1 - q_h) u\Big(w[(l - r_f)x + r_f] \Big) \Big], \\
J_l(w) &= \max_{0 \le x \le 1} \min_{q_l \in [\underline{q}^l, \overline{q}^l]} \Big[q_l u\Big(w[(h - r_f)x + r_f] \Big) + (1 - q_l) u\Big(w[(l - r_f)x + r_f] \Big) \Big], \\
J_0(w_0) &= \max_{0 \le x \le 1} \min_{q_0 \in [\underline{q}^0, \overline{q}^0]} \Big[q_0 J_h\Big(w_0[(h - r_f)x + r_f] \Big) + (1 - q_0) J_l\Big(w_0[(l - r_f)x + r_f] \Big) \Big].
\end{aligned}$$

Knox stresses that in each of the three subproblems, the investor behaves as if there were a separate asset whose return uncertainty is described by an interval of probabilities.

In our specification for $\mathcal{C}$ in (5), we have imposed exchangeability: for each $q \in \mathcal{C}$,

$$\Pr_q(R_0 = h, R_1 = l) = \Pr_q(R_0 = l, R_1 = h).$$

This restriction is not compatible with a rectangular set of priors, and so we may not want to impose the rectangularity restriction.

A related issue is discussed by Epstein & Schneider in the context of a dynamic, three-color Ellsberg urn experiment in which there are 30 balls with color A and 60 balls with color B or C. (In Epstein & Schneider, A = red, B = blue, C= green.) A ball is drawn at random from the urn at time 0. A bet on $(1, 0, 1)$ pays off one util if the color is A or C, and $(0, 1, 1)$ pays one util if the color is B or C. At $t = 1$ the decision-maker is told whether or not the color is C and is asked to choose between $(1, 0, 1)$ and $(0, 1, 1)$. The state space is $\{A, B, C\}$. A prior distribution consists of three probabilities (q_A, q_B, q_C) that are nonnegative and sum to 1. Consider the set of priors

$$\mathcal{C} = \{q = (\frac{1}{3}, q_B, \frac{2}{3} - q_B) : \frac{1}{6} \leq q_B \leq \frac{1}{2}\}.$$

The decision maker can form a contingent plan. For example, choose $(1, 0, 1)$ if C^c (not C) and choose x if C, where x could be $(1, 0, 1)$ or $(0, 1, 1)$. If C, then $(1, 0, 1)$ and $(0, 1, 1)$ both pay 1 util, and so, for either choice of x, the value of the plan is

$$\min_{q \in \mathcal{C}} [\Pr_q(C^c) \Pr_q(A \,|\, C^c) \cdot 1 + \Pr_q(C) \cdot 1] = \min_{q \in \mathcal{C}} [\Pr_q(A) + \Pr_q(C)] = \frac{1}{2}.$$

Likewise the plan that chooses $(0, 1, 1)$ if C^c has value

$$\min_{q \in \mathcal{C}} [\Pr_q(C^c) \Pr_q(B \,|\, C^c) \cdot 1 + \Pr_q(C) \cdot 1] = \min_{q \in \mathcal{C}} [\Pr_q(B) + \Pr_q(C)] = \frac{2}{3}.$$

So $(0, 1, 1)$ is chosen over $(1, 0, 1)$.

This choice is often regarded as intuitive, but the set of priors is not rectangular. Epstein & Schneider say that their modeling approach would suggest replacing $\mathcal{C}$ by the smallest rectangular set containing $\mathcal{C}$, which they give:

$$\mathcal{C}^{rect} = \Big\{ \Big(\frac{1}{3} \frac{\frac{1}{3} + q'_B}{\frac{1}{3} + q_B}, q_B \frac{\frac{1}{3} + q'_B}{\frac{1}{3} + q_B}, \frac{2}{3} - q'_B \Big) : \frac{1}{6} \leq q_B, q'_B \leq \frac{1}{2} \Big\}.$$

In the rectangular prior there is a range of probabilities for A, even though it is given that the fraction of A balls in the urn is 1/3. So the decision maker may not want to impose the rectangularity restriction. (With the rectangular set of priors, $(1, 0, 1)$ is chosen over $(0, 1, 1)$, reversing the choice based on the original set of priors $\mathcal{C}$.)

2.2.2. Conditional Preferences. Epstein & Schneider develop a conditional preference ordering conditional on the information available at each date. Imposing dynamic consistency across these preference orderings leads to the rectangularity restriction on the set of priors. Knox argues that the problematic aspects arise because these conditional preferences impose consequentialism, the property that counterfactuals are ignored. This is further developed in Hanany & Klibanoff (2007, p. 262):

> As Machina (1989) has emphasized, once we move beyond expected utility and preferences are not separable across events, updating in a dynamically

> consistent way entails respecting these non-separabilities by allowing updated preferences to depend on more than just the conditioning event. For this reason, we will see that dynamic consistency naturally leads a decision maker (DM) concerned with ambiguity to adopt rules for updating beliefs that depend on prior choices and/or the feasible set for the problem.

Hanany & Klibanoff also use a version of Ellsberg's three-color problem as a motivating example. The urn contains 90 balls, 30 of which are known to be A and 60 of which are somehow divided between B and C, with no further information on the distribution. (In Hanany & Klibanoff, A = black, B = red, C = yellow, and the urn contains 120 balls, with 1/3 A and 2/3 B or C.) A ball is to be drawn at random from the urn, and the DM faces a choice among bets paying off depending on the color of the drawn ball. Any such bet can be written as a triple $(u_A, u_B, u_C) \in \mathbb{R}^3$, where each ordinate represents the payoff if the respective color is drawn. Typical preferences have $(1, 0, 0) \succ (0, 1, 0)$ and $(0, 1, 1) \succ (1, 0, 1)$, reflecting a preference for the less ambiguous bets. In the dynamic version of the problem there is an interim stage where the DM is told whether or not the drawn ball is C. Two choice pairs are considered. In choice pair 1, if the drawn ball is C then the payoff is 0. If not C, so conditional on the event $E = \{A, B\}$, then the DM chooses between A and B. The choice "Bet on A" leads to the payoff vector $(1, 0, 0)$ whereas the choice "Bet on B" leads to payoffs $(0, 1, 0)$. In choice pair 2, if the drawn ball is C then the payoff is 1. If not C, then the DM chooses between A and B. Now the choice "Bet on A" leads to the payoff vector $(1, 0, 1)$ whereas the choice "Bet on B" leads to payoffs $(0, 1, 1)$. See Figure 3. Hanany & Klibanoff argue that in choosing between the bets $(1, 0, 0)$ and $(0, 1, 0)$, the opportunity to condition the choice on the information at the interim stage does not change the problem in an essential way. Therefore preference should remain $(1, 0, 0) \succ (0, 1, 0)$ and $(0, 1, 1) \succ (1, 0, 1)$, as in the original problem. Hanany & Klibanoff conclude that these preferences are inconsistent with backward induction, which requires the DM to snip the tree at the node following the event $\{A, B\}$ and to choose as if this were the entire problem.

> But then the choice between $(1, 0, 0)$ and $(0, 1, 0)$ must be the same as the choice between $(1, 0, 1)$ and $(0, 1, 1)$ since the snipped trees for the two choice pairs are identical, rendering the Ellsberg choices impossible. It follows that no model of dynamic choice under ambiguity implying backward induction can deliver the Ellsberg preferences in this example.

Hanany & Klibanoff (2007, p. 264).

This point is, I think, fundamental. The Ellsberg paradox has been a major motivation for developing models for preferences that distinguish between roulette lotteries and horse lotteries, allowing for ambiguity aversion. In the dynamic version of these preferences, the stress has been on recursive models, which can be solved by backward induction. Conditional preferences that have a recursive form are very convenient for computation, but there is a tension here with the motivating Ellsberg intuition in which conditional preferences are not recursive. If the goal is the positive one of modeling observed behavior, then recursive preferences may not be suitable. My goal is the normative one of adding robustness to SEU preferences, so the failure of recursive preferences to model dynamic Ellsberg behavior is less of a concern.

Hanany & Klibanoff develop conditional preferences in the maxmin expected utility

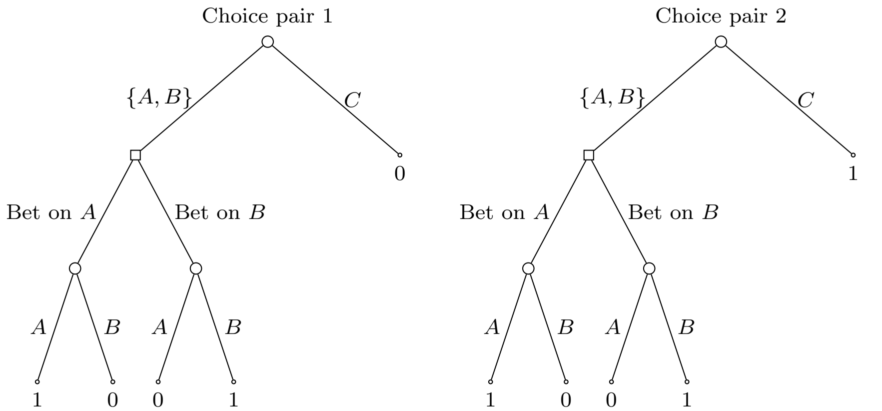

Figure 3

Dynamic Ellsburg problem with urn containing 30 A balls and 60 balls that are either B or C. The decision maker is told whether the drawn ball is C before betting on A or B. The payoffs for A and B are 1. The payoff for C is 0 in choice problem 1, and the payoff is 1 in choice problem 2.

(MEU) framework. These preferences are dynamically consistent in that ex ante optimal contingent choices are respected when a planned-for contingency arises. Their general results provide update rules for MEU preferences that apply Bayes' rule to some of the probability measures used in representing the DM's unconditional preferences. They apply their general results to the dynamic Ellsberg problem. In their setup, for any MEU preferences over payoff vectors in $\mathbb{R}^3$, there exists a convex set of probability measures, $\mathcal{C}$, over the three colors and a utility function, $u\colon \mathbb{R} \to \mathbb{R}$, such that for all $f, g \in \mathbb{R}^3$, $f \succeq g \Longleftrightarrow \min_{q\in\mathcal{C}} \int (u \circ f)\, dq \geq \min_{q\in\mathcal{C}} \int (u \circ g)\, dq$. Let $u(x) = x$ for all $x \in \mathbb{R}$, and let $\mathcal{C} = \{(\frac{1}{3}, \alpha, \frac{2}{3} - \alpha : \alpha \in [\frac{1}{4}, \frac{5}{12}]\}$, a set of measures symmetric with respect to the probabilities of B and C. According to these preferences, $(1,0,0) \succ (0,1,0)$ and $(0,1,1) \succ (1,0,1)$. They show that dynamically consistent updating in the Ellsberg problem corresponds to updating the set of measures to be any closed, convex subset of $\mathcal{C}^1_E = \{(\alpha, 1-\alpha, 0) : \alpha \in [\frac{1}{2}, \frac{4}{7}]\}$ in choice problem 1, and any closed, convex subset of $\mathcal{C}^2_E = \{(\alpha, 1-\alpha, 0) : \alpha \in [\frac{4}{9}, \frac{1}{2}]\}$ in choice problem 2.

This result corresponds to updating the least-favorable priors in the unconditional problem. Let $f = (1,0,0)$ and $g = (0,1,0)$. The least-favorable prior q^* satisfies

$$q^* = \arg\min_{q\in\mathcal{C}} \Big[\max\{\int (u \circ f)\, dq, \int (u \circ g)\, dq\}\Big].$$

Note that

$$\max\{\int (u \circ f)\, dq, \int (u \circ g)\, dq\} = \begin{cases} \frac{1}{3}, & \text{if } q(B) \leq \frac{1}{3}; \\ q(B), & \text{otherwise.} \end{cases}$$

The minimum over $q \in \mathcal{C}$ is achieved by any $q^* = (\frac{1}{3}, \alpha, \frac{2}{3} - \alpha)$ with $\alpha \in [\frac{1}{4}, \frac{5}{12}]$ and $\alpha \leq \frac{1}{3}$. So the set of least-favorable priors is

$$Q = \{q^* = (\frac{1}{3}, \alpha, \frac{2}{3} - \alpha : \alpha \in [\frac{1}{4}, \frac{1}{3}]\}.$$

Now condition on the event $E = \{A, B\}$ and update a least-favorable prior using Bayes' rule:

$$q^*(A \mid E) = \frac{1}{1 + 3q^*(B)} \in [\frac{1}{2}, \frac{4}{7}].$$

So updating the set of least-favorable priors gives

$$Q_E = \{(\alpha, 1-\alpha, 0) : \alpha \in [\frac{1}{2}, \frac{4}{7}]\} = \mathcal{C}^1_E.$$

The conditional preferences have $(1,0,0) \succ (0,1,0)$:

$$\min_{q \in \mathcal{C}^1_E} \int (u \circ f)\, dq = \min_{q \in \mathcal{C}^1_E} q(A) = \frac{1}{2} > \min_{q \in \mathcal{C}^1_E} \int (u \circ g)\, dq = \min_{q \in \mathcal{C}^1_E} q(B) = \frac{3}{7}.$$

In choice problem 2, let $f = (1,0,1)$ and $g = (0,1,1)$. Updating the set of least-favorable priors gives $\mathcal{C}^2_E$. The conditional preferences have $(1,0,1) \prec (0,1,1)$:

$$\min_{q \in \mathcal{C}^2_E} \int (u \circ f)\, dq = \min_{q \in \mathcal{C}^2_E} q(A) = \frac{4}{9} < \min_{q \in \mathcal{C}^2_E} \int (u \circ g)\, dq = \min_{q \in \mathcal{C}^2_E} q(B) = \frac{1}{2}.$$

2.3. Multiplier Preferences

I follow Strzalecki (2011) in setting up multiplier preferences. They are based on a reference probability model q. Other probability models p are considered, but they are penalized by the relative entropy $R(\cdot \,\|\, q)$, which is a mapping from $\triangle(S)$, the set of probability distributions on the state space S, into $[0, \infty]$:

$$R(p \,\|\, q) = \begin{cases} \int_S \left(\log \frac{dp}{dq}\right) dp, & \text{if } p \text{ is absolutely continuous with respect to } q; \\ \infty, & \text{otherwise.} \end{cases} \tag{6}$$

The set Z denotes the possible consequences and $\triangle(Z)$ denotes probability distributions on Z with finite support. Let Σ denote a sigma-algebra of events in S. An act f is a finite-valued Σ-measurable mapping from the state space S to lotteries over consequences: $f \colon S \to \triangle(Z)$; the set of all such acts is denoted $\mathcal{F}(\triangle(Z))$. Acts f are ranked according to the criterion

$$V(f) = \min_{p \in \triangle(S)} \int_S u(f(s))\, dp(s) + \kappa R(p \,\|\, q), \tag{7}$$

where $u \colon \triangle(Z) \to \mathbb{R}$ is a nonconstant, affine function, $\kappa \in (0, \infty)$, and $q \in \triangle(S)$. Define a class of transformations $\zeta_\kappa \colon \mathbb{R} \to \mathbb{R}$ which are strictly increasing and concave:

$$\zeta_\kappa(u) = -\exp(-\frac{u}{\kappa}), \tag{8}$$

with $\zeta_\kappa^{-1}(u) = -\kappa \log(-u)$. There is a very useful variational formula in Proposition 1.4.2 of Dupuis & Ellis (1997): for any bounded, measurable function $h \colon S \to \mathbb{R}$ and $q \in \triangle(S)$,

$$\min_{p \in \triangle(S)} \int_S h(s)\, dp(s) + \kappa R(p \,\|\, q) = \zeta_\kappa^{-1}\Big(\int_S (\zeta_\kappa \circ h)(s)\, dq(s)\Big).$$

The minimum is attained uniquely at p^*, which has the following density with respect to q:

$$\frac{dp^*}{dq}(s) = \frac{\exp(-\frac{1}{\kappa} h(s))}{\int_S \exp(-\frac{1}{\kappa} h(s))\, dq(s)}.$$

Hence

$$V(f) = -\kappa \log\Big(\int_S \exp(-\frac{1}{\kappa}u(f(s)))\,dq(s)\Big), \quad \frac{dp^*}{dq}(s) = \frac{\exp(-\frac{1}{\kappa}u(f(s)))}{\int_S \exp(-\frac{1}{\kappa}u(f(s)))\,dq(s)}. \tag{9}$$

In our two-period portfolio choice problem, let the state space be

$$S = \{(h,h),(h,l),(l,h),(l,l)\} \tag{10}$$

corresponding to the possible values for the returns (R_0, R_1). The investor has a contingent plan in which the fraction of wealth invested in the risky asset at $t = 0$ is a_0; at $t = 1$, it is $a_1(h)$ if $R_0 = h$ and $a_1(l)$ if $R_0 = l$. The payoff profile f for this plan is given by

$$\begin{aligned} f(h,h) &= w_0[(h-r_f)a_0 + r_f][(h-r_f)a_1(h) + r_f], \\ f(h,l) &= w_0[(h-r_f)a_0 + r_f][(l-r_f)a_1(h) + r_f], \\ f(l,h) &= w_0[(l-r_f)a_0 + r_f][(h-r_f)a_1(l) + r_f], \\ f(l,l) &= w_0[(l-r_f)a_0 + r_f][(l-r_f)a_1(l) + r_f]. \end{aligned} \tag{11}$$

For the reference probabilty model q, we can use the predictive distribution in (1):

$$\begin{aligned} \Pr_q(R_0 = h) &= \frac{\alpha}{\alpha+\beta} = q_0, \\ \Pr_q(R_1 = h \,|\, R_0 = h) &= \frac{\alpha+1}{\alpha+\beta+1} = q_h, \\ \Pr_q(R_1 = h \,|\, R_0 = l) &= \frac{\alpha}{\alpha+\beta+1} = q_l. \end{aligned}$$

Note that this predictive distribution is exchangeable:

$$\Pr_q(R_0 = h, R_1 = l) = \Pr_q(R_0 = l, R_1 = h) = \frac{\alpha\beta}{(\alpha+\beta)(\alpha+\beta+1)}.$$

Given $p \in \triangle(S)$, use the notation

$$p_0 = \Pr_p(R_0 = h),\;\; p_h = \Pr_p(R_1 = h \,|\, R_0 = h),\;\; p_l = \Pr_p(R_1 = h \,|\, R_0 = l).$$

The investor's problem is

$$\max_{0 \le a \le 1} \min_{p \in \triangle(S)} \Big(\int_S u(f(s))\,dp(s) + \kappa R(p \,\|\, q)\Big). \tag{12}$$

Note that

$$\min_{p \in \triangle(S)} \int_S u(f(s))\,dp(s) + \kappa R(p \,\|\, q) = -\kappa \log\Big(-\int_S \zeta_\kappa(u(f(s)))\,dq(s)\Big).$$

So we can replace the utility function u by $\zeta_\kappa \circ u$. The function $\zeta_\kappa \circ u$ is not a von Neumann-Morgenstern (vNM) utility function—it is not affine on $\Delta(Z)$. This does not matter because, for each $s \in S$, the distribution of $f(s)$ assigns probability one to a single point. So we can apply SEU preferences using $\zeta_\kappa \circ u$ and the reference distribution q. This can be solved by dynamic programming, as in Section 2.1. Define the maximized value functions

$$\begin{aligned} J_h(w) &= \max_{0 \le x \le 1} \Big[q_h \zeta_\kappa \circ u\Big(w[(h-r_f)x + r_f]\Big) + (1-q_h)\zeta_\kappa \circ u\Big(w[(l-r_f)x + r_f]\Big)\Big], \\ J_l(w) &= \max_{0 \le x \le 1} \Big[q_l \zeta_\kappa \circ u\Big(w[(h-r_f)x + r_f]\Big) + (1-q_l)\zeta_\kappa \circ u\Big(w[(l-r_f)x + r_f]\Big)\Big]. \end{aligned}$$

Then the optimal portfolio weight on the risky asset in period 0 is

$$a_0^* = \arg \max_{0 \le x \le 1} \Big[q_0 J_h \Big(w_0[(h - r_f)x + r_f] \Big) + (1 - q_0) J_l \Big(w_0[(l - r_f)x + r_f] \Big) \Big]. \tag{13}$$

In our application of multiplier preferences, the reference distribution q for (R_0, R_1) is exchangeable, but the alternative distributions p are not constrained to be exchangeable. We can impose this restriction by working with a different state space. Now let $S = [0, 1]$. Conditional on $\theta \in S$, the distribution of (R_0, R_1) is

$$\begin{aligned} \Pr_\theta(R_0 = h, R_1 = h) &= \theta^2, \\ \Pr_\theta(R_0 = h, R_1 = l) &= \theta(1 - \theta), \\ \Pr_\theta(R_0 = l, R_1 = h) &= (1 - \theta)\theta, \\ \Pr_\theta(R_0 = l, R_1 = l) &= (1 - \theta)^2. \end{aligned} \tag{14}$$

The investor has a contingent plan in which the fraction of wealth invested in the risky asset at $t = 0$ is a_0; at $t = 1$, it is $a_1(h)$ if $R_0 = h$ and $a_1(l)$ if $R_0 = l$. The payoff profile g for this plan maps states $\theta \in S$ into lotteries over consequences. In state θ, the lottery assigns probabilities θ^2, $\theta(1-\theta)$, $(1-\theta)\theta$, $(1-\theta)^2$ to the consequences $f(h, h)$, $f(h, l)$, $f(l, h)$, $f(l, l)$ (where f is defined in (11)). So

$$\begin{aligned} u(g(\theta)) = {} & [\theta^2 f(h, h)^{1-\gamma} + \theta(1 - \theta) f(h, l)^{1-\gamma} \\ & + (1 - \theta)\theta f(l, h)^{1-\gamma} + (1 - \theta)^2 f(l, l)^{1-\gamma}]/(1 - \gamma). \end{aligned}$$

For the reference probability model q, we shall use a beta distribution with parameters α and β.

The investors's problem is

$$\begin{aligned} & \max_{0 \le a \le 1} \min_{p \in \triangle([0,1])} \Big(\int_0^1 u(g(\theta))\, dp(\theta) + \kappa R(p \,\|\, q) \Big) \\ & = \max_{0 \le a \le 1} -\kappa \log \Big(- \int_0^1 \zeta_\kappa(u(g(\theta)))\, dq(\theta) \Big). \end{aligned}$$

Now it matters that $\zeta_\kappa \circ u$ is not a vNM utility function. The optimal portfolio weights can be obtained by maximizing the following objective function with respect to a_0, $a_1(h)$, $a_1(l)$:

$$\begin{aligned} & -\kappa \log \int_0^1 \exp\Big(-\frac{1}{\kappa} \frac{w_0^{1-\gamma}}{(1 - \gamma)} \Big(\theta^2 \big[[(h - r_f)a_0 + r_f][(h - r_f)a_1(h) + r_f] \big]^{1-\gamma} \\ & \quad + \theta(1 - \theta) \big[[(h - r_f)a_0 + r_f][(l - r_f)a_1(h) + r_f] \big]^{1-\gamma} \\ & \quad + (1 - \theta)\theta \big[[(l - r_f)a_0 + r_f][(h - r_f)a_1(l) + r_f] \big]^{1-\gamma} \\ & \quad + (1 - \theta)^2 \big[[(l - r_f)a_0 + r_f][(l - r_f)a_1(l) + r_f] \big]^{1-\gamma} \Big) \Big) \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha-1}(1 - \theta)^{\beta-1}\, d\theta. \end{aligned} \tag{15}$$

When we impose the exchangeability restriction (on the distribution of (R_0, R_1)), we lose the recursive structure of preferences and the ability to solve the problem by dynamic programming. A similar issue arises in the dynamic Ellsberg three-color problem in Figure 3. Suppose we are given that one third of the balls in the urn are A, with two thirds either B or C. The decision maker is told whether the ball drawn is C before betting on A or B.

In choice pair 1 the payoff for C equals 0, and the choice "Bet on A" leads to the payoff vector $(1,0,0)$ whereas the choice "Bet on B" leads to payoffs $(0,1,0)$. In choice pair 2 the payoff for C equals 1, and the choice "Bet on A" leads to the payoff $(1,0,1)$ whereas the choice "Bet on B" leads to the payoff $(0,1,1)$. In setting up the problem using multiplier preferences, the reference probability model would assign probability 1/3 to drawing A. In general we would also want the alternative distribution p to assign probability 1/3 to drawing A, but consider setting up the problem without imposing that restriction. The state space is $S = \{A, B, C\}$ and $\triangle$ is the unit simplex in $\mathbb{R}^3$. The decision maker forms a contingent plan. For example in problem 1, choose $(1,0,0)$ if not C and choose x if C, where x could be $(1,0,0)$ or $(0,1,0)$. If C, then $(1,0,0)$ and $(0,1,0)$ both have a payoff of zero, and so, for either choice of x, the value of the plan is

$$\begin{aligned}
&\min_{p\in\triangle}\big[\mathrm{Pr}_p(C^c)\mathrm{Pr}_p(A \mid C^c)\cdot 1 + \mathrm{Pr}_p(C)\cdot 0 + \kappa R(p \,\|\, q)\big] \\
&= \min_{p\in\triangle}\big[\mathrm{Pr}_p(A)\cdot 1 + \mathrm{Pr}_p(C)\cdot 0 + \kappa R(p \,\|\, q)\big] \\
&= -\kappa \log\Big[\mathrm{Pr}_q(A)\exp(-\frac{1}{\kappa}\cdot 1) + \mathrm{Pr}_q(B)\exp(-\frac{1}{\kappa}\cdot 0) + \mathrm{Pr}_q(C)\exp(-\frac{1}{\kappa}\cdot 0)\Big].
\end{aligned}$$

Likewise the plan that chooses $(0,1,0)$ has value

$$-\kappa \log\Big[\mathrm{Pr}_q(A)\exp(-\frac{1}{\kappa}\cdot 0) + \mathrm{Pr}_q(B)\exp(-\frac{1}{\kappa}\cdot 1) + \mathrm{Pr}_q(C)\exp(-\frac{1}{\kappa}\cdot 0)\Big].$$

In problem 2, the plan that chooses $(1,0,1)$ has value

$$-\kappa \log\Big[\mathrm{Pr}_q(A)\exp(-\frac{1}{\kappa}\cdot 1) + \mathrm{Pr}_q(B)\exp(-\frac{1}{\kappa}\cdot 0) + \mathrm{Pr}_q(C)\exp(-\frac{1}{\kappa}\cdot 1)\Big],$$

and the plan that chooses $(0,1,1)$ has value

$$-\kappa \log\Big[\mathrm{Pr}_q(A)\exp(-\frac{1}{\kappa}\cdot 0) + \mathrm{Pr}_q(B)\exp(-\frac{1}{\kappa}\cdot 1) + \mathrm{Pr}_q(C)\exp(-\frac{1}{\kappa}\cdot 1)\Big].$$

So the decison maker makes the same choice on A versus B in problems 1 and 2, choosing A if $\mathrm{Pr}_q(A) > \mathrm{Pr}_q(B)$. We can obtain these solutions by snipping the trees at the decision nodes following the event $E = \{A, B\}$, setting the state space to $\{A, B\}$, and applying multiplier preferences using q' as the reference distribution, where q' is the Bayesian update of q:

$$\mathrm{Pr}_{q'}(A) = 1 - \mathrm{Pr}_{q'}(B) = \mathrm{Pr}_q(A \mid E) = \mathrm{Pr}_q(A)/\big(\mathrm{Pr}_q(A) + \mathrm{Pr}_q(B)\big).$$

If $\mathrm{Pr}_q(A) = \mathrm{Pr}_q(B) = \mathrm{Pr}_q(C) = 1/3$, then $(1,0,0) \sim (0,1,0)$ and $(1,0,1) \sim (0,1,1)$.

Now consider restricting $\mathrm{Pr}(A) = 1/3$ under p and q. The state space is $S = \{\theta : \theta \in [0,1]\}$ with $\mathrm{Pr}_\theta(B) = \frac{2}{3} - \mathrm{Pr}_\theta(C) = \frac{2}{3}\theta$. Suppose that

$$\int_0^1 \frac{2}{3}\theta \, dq(\theta) = \frac{1}{3}.$$

In problem 1 the contingent plan with "Bet on A" has value 1/3. The plan "Bet on B" has value

$$\min_{p\in\triangle([0,1])} \int_0^1 \frac{2}{3}\theta \, dp(\theta) + \kappa R(p \,\|\, q) = -\kappa \log\Big(\int_0^1 \exp(-\frac{1}{\kappa}\frac{2}{3}\theta)\, dq(\theta)\Big).$$

By Jensen's inequality,

$$\int_0^1 \exp(-\frac{1}{\kappa}\frac{2}{3}\theta)\, dq(\theta) > \exp\Big(-\frac{1}{\kappa}\frac{2}{3}\int_0^1 \theta\, dq(\theta)\Big) = \exp(-\frac{1}{\kappa}\frac{1}{3}),$$

and so $(1,0,0) \succ (0,1,0)$.

In problem 2 the contingent plan "Bet on B" has value 2/3. The plan "Bet on A" has value

$$\frac{1}{3} + \min_{p\in\triangle([0,1])} \int_0^1 \frac{2}{3}(1-\theta)\, dp(\theta) + \kappa R(p\,\|\,q) = \frac{1}{3} - \kappa \log\Big(\int_0^1 \exp\big(-\frac{1}{\kappa}\frac{2}{3}(1-\theta)\big)\, dq(\theta)\Big).$$

By Jensen's inequality this is less than $\frac{2}{3}$, and so $(0,1,1) \succ (1,0,1)$. So when we restrict $\Pr(A) = 1/3$ under p and q, we do not have the consequentialist solution that snips the trees at the decision nodes following the event $E = \{A, B\}$. Now our solution exhibits typical Ellsberg behavior.

Hansen & Miao (2018) explore the relative entropy relations between priors, likelihoods, and predictive densities in a static setting. There is a prior distribution π for parameter values $\theta \in \Theta$ and a likelihood λ for the density given θ for possible outcomes $y \in \mathcal{Y}$ (with respect to a measure τ). The predictive density for y is

$$\phi(y) = \int_\Theta \lambda(y\,|\,\theta)\,\pi(d\theta).$$

The reference distribution counterparts are $\hat\pi$, $\hat\lambda$, and $\hat\phi$. Hansen & Miao pose and solve two problems that adjust for robustness. First, robust evaluation of a y-dependent utility $U(y)$ gives

$$\begin{aligned}\min_{\phi} \int_{\mathcal{Y}} U(y)\phi(y)\tau(dy) + \kappa \int_{\mathcal{Y}} \Big[\log\phi(y) - \log\hat\phi(y)\Big]\phi(y)\tau(dy) \\ = -\kappa \log \int_{\mathcal{Y}} \exp\Big[-\frac{1}{\kappa}U(y)\Big]\hat\phi(y)\tau(dy).\end{aligned} \tag{16}$$

Second, they target prior robustness by restricting $\lambda = \hat\lambda$, eliminating specification concerns about the likelihood, and solving

$$\begin{aligned}\min_{\pi\in\Pi} \int_\Theta \overline{U}(\theta)\pi(d\theta) + \kappa \int_\Theta \log\Big[\frac{d\pi}{d\hat\pi}(\theta)\Big]\pi(d\theta) \\ = -\kappa\log\int_\Theta \exp\Big[-\frac{1}{\kappa}\overline{U}(\theta)\Big]\hat\pi(d\theta),\end{aligned} \tag{17}$$

where Π is the set of priors that are absolutely continuous with respect to $\hat\pi$ and

$$\overline{U}(\theta) \equiv \int_{\mathcal{Y}} U(y)\hat\lambda(y\,|\,\theta)\tau(dy).$$

The first problem corresponds to our first application of multiplier preferences to the portfolio choice problem, in which we did not restrict the set of distributions for (R_0, R_1). The second problem corresponds to our second application of multiplier preferences, in which we imposed the restriction that R_0 and R_1 are independent and identically distributed conditional on θ, where the marginal distribution of θ is unrestricted. Likewise, the first problem corresponds to our first application of multiplier preferences in the dynamic Ellsberg three-color problem, in which we did not restrict the alternative distribution p to assign probability 1/3 to drawing A. The second problem corresponds to imposing the restriction that $\Pr_p(A) = 1/3$.

2.4. Divergence Preferences and Constraint Preferences

Maccheroni et al. (2006a) weaken the Anscombe & Aumann independence axiom to weak certainty independence. Using the notation from Section 2.3, this says that if $f, g \in \mathcal{F}(\triangle(Z))$, $x, y \in \triangle(Z)$, and $\alpha \in (0,1)$, then

$$\begin{aligned} &\alpha f + (1-\alpha)x \succsim \alpha g + (1-\alpha)x \\ &\quad \Rightarrow \quad \alpha f + (1-\alpha)y \succsim \alpha g + (1-\alpha)y. \end{aligned}$$

This leads to variational preferences in which acts f are ranked according to the criterion

$$V(f) = \min_{p \in \triangle(S)} \int_S u(f(s))\, dp(s) + c(p), \tag{18}$$

where $u\colon \triangle(Z) \to \mathbb{R}$ is a nonconstant, affine function, and the cost function $c\colon \triangle(S) \to [0,\infty]$ is convex. This includes multiplier preferences as a special case with $c(p) = \kappa R(p \,\|\, q)$. It also includes MEU preferences as a special case with

$$c(p) = \begin{cases} 0, & \text{if } p \in \mathcal{C}; \\ \infty, & \text{otherwise.} \end{cases}$$

Strzalecki (2011) shows that multiplier preferences can be characterized by adding Savage's P2 axiom (applied to all Anscombe & Aumann acts) to the axioms of Maccheroni et al. With $E \in \Sigma$, let $f_E g$ denote an act with $f_E g(s) = f(s)$ if $s \in E$ and $g(s)$ if $s \notin E$. Then Strzalecki adds Axiom P2—Savage's Sure-Thing Principle: for all $E \in \Sigma$ and $f, g, h, h' \in \mathcal{F}(\triangle(Z))$,

$$f_E h \succsim g_E h \quad \Rightarrow \quad f_E h' \succsim g_E h'.$$

Consider an individual's conditional preferences upon learning that an event E has occurred. Following Machina & Schmeidler (1992), let an act h defined over the complement of E provide a counterfactual, and define a conditional of $\succsim$ on E given h (denoted by $\succsim_{E,h}$) by

$$f \succsim_{E,h} g \quad \text{if and only if} \quad fEh \succsim gEh.$$

Gumen & Savochkin (2013) note that dynamic consistency is used here as a definition of conditional preferences, rather than as a link between two exogenously given preference relations. Consequentialism can be imposed on preferences by requiring that all conditionals $\succsim_{E,h}$ be independent of h. In this setting, this is equivalent to Savage's Sure-Thing Principle. Gumen & Savochkin show that multiplier preferences are stable in that the corresponding conditional preferences $\succsim_{E,h}$ (which are independent of h) belong to the class of multiplier preferences.

Divergence preferences are an important class of variational preferences, and they include multiplier preferences as a special case. There is an underlying probability measure $q \in \triangle(S)$. Given a a convex, continuous function $\phi : \mathbb{R}_+ \to \mathbb{R}$ such that $\phi(1) = 0$ and $\lim_{t\to\infty} \phi(t)/t = \infty$, the ϕ-divergence of $p \in \triangle(S)$ with respect to q is given by

$$D_\phi(p \,\|\, q) = \begin{cases} \int_S \phi\left(\frac{dp}{dq}\right) dq, & \text{if } p \text{ is absolutely continuous with respect to } q; \\ \infty, & \text{otherwise.} \end{cases}$$

Acts are ranked according to the criterion

$$V(f) = \min_{p \in \triangle(S)} \int_S u(f(s))\, dp(s) + \kappa D_\phi(p \,\|\, q), \tag{19}$$

where $\kappa \in (0, \infty)$. Setting $\phi(t) = t\log(t)$, these divergence preferences coincide with multiplier preferences in (7). Maccheroni et. al (2006a) also consider $\phi(t) = 2^{-1}(t-1)^2$ (χ^2-divergence). They refer to these divergence preferences as Gini preferences and relate them to mean-variance preferences. Other examples of ϕ-divergence appear in the literature on generalized empirical likelihood (Newey & Smith (2004)); in Kitamura et al. (2013) on robust estimation using neighborhoods based on squared Hellinger distance; and in Christensen & Connault (2019) on global sensitivity analysis using convex duality to obtain dual representations of infinite-dimension optimization problems as low-dimension convex programs.

The optimization problem in (19) can be simplified using convex duality. Let $h(s) = \frac{dp}{dq}(s)$. Then

$$V(f) = \inf_{h(\,\cdot\,)\geq 0, \int h\, dq = 1} \int_S u(f(s))h(s)\, dq(s) + \kappa \int_S \phi(h(s))\, dq(s).$$

Setting up the Lagrangian,

$$V(f) = \max_{\delta\in\mathbb{R}} \left(\inf_{h(\,\cdot\,)\geq 0} \int_S [u(f(s))h(s) + \kappa\phi(h(s)) + \delta h(s)]\, dq(s) - \delta \right)$$

(Luenberger (1969), Chap. 8.6, Theorem 1 and Chap. 8.8, Problem 7). Note that

$$\begin{aligned} &\inf_{h(\,\cdot\,)\geq 0} \int_S [u(f(s))h(s) + \kappa\phi(h(s)) + \delta h(s)]\, dq(s) \\ &\quad \geq \int_S \inf_{t\in\mathbb{R}_+} [u(f(s))t + \kappa\phi(t) + \delta t]\, dq(s), \end{aligned} \tag{20}$$

and

$$\inf_{t\in\mathbb{R}_+} [u(f(s))t + \kappa\phi(t) + \delta t] = -\kappa \sup_{t\geq 0} [-\frac{u(f(s)) + \delta}{\kappa} t - \phi(t)] = -\kappa\phi^*(-\frac{u(f(s)) + \delta}{\kappa}),$$

where ϕ^* is the convex conjugate of ϕ with $\phi^* : \mathbb{R} \to (-\infty, \infty]$,

$$\phi^*(x) = \sup_{t\geq 0} [xt - \phi(t)].$$

The lower bound in (20) is achieved by setting

$$h(s) = \arg\min_{t\geq 0} [u(f(s))t + \kappa\phi(t) + \delta t] \tag{21}$$

(assuming the minimum in (21) is attained), and then

$$V(f) = \max_{\delta\in\mathbb{R}} -\kappa \int_S \phi^*(-\frac{u(f(s)) + \delta}{\kappa})\, dq(s) - \delta. \tag{22}$$

For example, with $\phi(t) = t\log t$ and $\phi^*(x) = \exp(x-1)$, the simplified expression for divergence preferences $V(f)$ in (22) coincides with the simplified expression for multiplier preferences $V(f)$ in (9).

Consider the investor's problem in (12), but replacing relative entropy $R(p\,\|\,q)$ with ϕ-divergence $D_\phi(p\,\|\,q)$:

$$\max_{0\leq a\leq 1} \min_{p\in\triangle(S)} \Big(\int_S u(f(s))\, dp(s) + \kappa D_\phi(p\,\|\,q) \Big). \tag{23}$$

The state space as in (10) is $S = \{(h,h),(h,l),(l,h),(l,l)\}$, and the act f is in (11). Note that

$$\begin{aligned}
&\max_{0\le a\le 1}\min_{p\in\triangle(S)}\Big(\int_S u(f(s))\,dp(s) + \kappa D_\phi(p\,\|\,q)\Big)\\
&\quad= \max_{0\le a\le 1}\max_{\delta\in\mathbb{R}}\Big(-\kappa\int_S \phi^*(-\frac{u(f(s))+\delta}{\kappa})\,dq(s) - \delta\Big)\\
&\quad= \max_{\delta\in\mathbb{R}}\Big(\max_{0\le a\le 1}\Big(-\kappa\int_S \phi^*(-\frac{u(f(s))+\delta}{\kappa})\,dq(s)\Big) - \delta\Big). \qquad (24)
\end{aligned}$$

For a given value of δ, define the increasing, concave function $b_\delta : \mathbb{R} \to [-\infty,\infty)$ by

$$b_\delta(x) = -\kappa\phi^*(-\frac{x+\delta}{\kappa}).$$

Replace the utility function u by $b_\delta \circ u$ and solve

$$\max_{0\le a\le 1}\int_S b_\delta \circ u(f(s))\,dq(s). \qquad (25)$$

The function $b_\delta \circ u$ is not a vNM utility function—it is not affine on $\Delta(Z)$. This does not matter because, for each $s \in S$, the distribution of $f(s)$ assigns probability one to a single point. So we can apply SEU preferences using $b_\delta \circ u$ and the reference distribution q. Then (25) can be solved by dynamic programming, as in Section 2.1 and (13). The solution in (24) is completed by maximizing over δ in $\mathbb{R}$.

Hansen & Sargent (2001) and Hansen et al. (2006) relate multiplier preferences to the MEU preferences of Gilboa & Schmeidler (1989). Acts f are ranked according to the criterion

$$V(f) = \min_{p\in\triangle(S):R(p\,\|\,q)\le\eta}\int_S u(f(s))\,dp(s), \qquad (26)$$

where $\eta > 0$ fixes an entropy neighborhood of the reference distribution q, providing the set $\mathcal{C}$ of priors in MEU. Hansen & Sargent refer to (26) as constraint preferences. Multiplier preferences in (7) and constraint preferences in (26) are distinct, but they generate the same optimal portfolio strategies for corresponding values of η and κ . This holds more generally, connecting divergence preferences in (19) and constraint preferences defined by

$$V(f) = \min_{p\in\triangle(S):D_\phi(p\,\|\,q)\le\eta}\int_S u(f(s))\,dp(s). \qquad (27)$$

We can see this by applying a minimax theorem. For example, suppose that the set of feasible portfolio strategies is finite: $a = (a_0, a_1(l), a_1(h)) \in A$ with the finite set $A = \{a^{(j)}\}_{j=1}^J$, and we allow mixed strategies. The probability weights in the mixed strategy are $\alpha \in \Lambda$, with $\Lambda = \{(\alpha_1,\ldots,\alpha_J) : \alpha_j \ge 0, \sum_{j=1}^J \alpha_j = 1\}$. Let $f^{(j)}$ denote the act corresponding to $a^{(j)}$ and define

$$W(\alpha,p) = \sum_{j=1}^J \alpha_j \int_S u(f^{(j)}(s))\,dp(s).$$

Then the minimax theorem for S games implies that

$$\max_{\alpha\in\Lambda}\min_{p\in\triangle(S):D_\phi(p\,\|\,q)\le\eta} W(\alpha,p) = \min_{p\in\triangle(S):D_\phi(p\,\|\,q)\le\eta}\max_{\alpha\in\Lambda} W(\alpha,p), \qquad (28)$$

and

$$\max_{\alpha\in\Lambda}\min_{p\in\triangle(S)} W(\alpha,p)+\kappa D_\phi(p\,\|\,q)$$
$$=\min_{p\in\triangle(S)}\max_{\alpha\in\Lambda} W(\alpha,p)+\kappa D_\phi(p\,\|\,q). \tag{29}$$

(Blackwell & Girshick (1954), Theorem 2.4.2 and Ferguson (1967), Chap. 2, Theorem 1.) Fix a value for $\kappa>0$ and let (α^*,p^*) denote a solution to the multiplier problem in (29):

$$\min_{p\in\triangle(S)}\max_{\alpha\in\Lambda} W(\alpha,p)+\kappa D_\phi(p\,\|\,q)$$
$$=\max_{\alpha\in\Lambda} W(\alpha,p^*)+\kappa D_\phi(p^*\,\|\,q)=W(\alpha^*,p^*)+\kappa D_\phi(p^*\,\|\,q).$$

Then, with $\eta=D_\phi(p^*\,\|\,q)$, (α^*,p^*) solves the constraint problem in (28):

$$\min_{p\in\triangle(S):D_\phi(p\,\|\,q)\le\eta}\max_{\alpha\in\Lambda} W(\alpha,p)=\max_{\alpha\in\Lambda} W(\alpha,p^*)=W(\alpha^*,p^*)$$

(Luenberger (1969), Chap. 8.4, Theorem 1). So we can interpret p^* as a least-favorable prior for the corresponding constraint preferences with $\eta=D_\phi(p^*\,\|\,q)$.

If we focus on constraint preferences, then the connection with divergence preferences can provide an algorithm. Solve the portfolio choice problem using divergence preferences with a range of values for the multiplier κ (taking advantage of the simplified form for $V(f)$ in (22) or, with multiplier preferences, in (9)). This will provide solutions to the portfolio choice problem using constraint preferences for a range of values of the constraint parameter η. Then the choice of η will fix a convex set $\mathcal{C}$ of priors for MEU preferences. Good (1952) considers a set of priors and argues that a minimax solution is reasonable provided that only reasonable subjective distributions are entertained. (See also Chamberlain (2000a).) So the choice of the constraint parameter η would be based on considering the set of implied predictive distributions for portfolio returns. Connecting optimal portfolios for divergence preferences and constraint preferences is useful because the optimal portfolio problem with divergence preferences can be solved using dynamic programming (as in (25) and (24)), when we do not restrict the set of distributions for (R_0,R_1) by imposing the reference likelihood function. The solution simplifies using multiplier preferences (as in (13)).

We can also work directly with constraint preferences. Setting up the Lagrangian,

$$\begin{aligned}V(f)&=\min_{p\in\triangle(S):D_\phi(p\,\|\,q)\le\eta}\int_S u(f(s))\,dp(s)\\&=\max_{\delta_1\in\mathbb{R},\delta_2\in\mathbb{R}_+}\Big(\inf_{h(\cdot)\ge 0}\int_S[u(f(s))h(s)+\delta_1 h(s)+\delta_2\phi(h(s))]\,dq(s)-\delta_1-\delta_2\eta\Big)\\&=\max_{\delta_1\in\mathbb{R},\delta_2\in\mathbb{R}_+}\Big(-\int(\delta_2\phi)^*(-(u(f(s))+\delta_1))\,dq(s)-\delta_1-\delta_2\eta\Big),\end{aligned}$$

where $(\delta_2\phi)^*$ is the convex dual of $\delta_2\phi$:

$$(\delta_2\phi)^*(x)=\begin{cases}\delta_2\phi^*(\frac{x}{\delta_2}), & \text{if } \delta_2>0;\\ \sup_{t\ge 0}xt, & \text{if } \delta_2=0.\end{cases}$$

For a given value of $\delta=(\delta_1,\delta_2)$, define

$$b_\delta(x)=-(\delta_2\phi)^*(-(x+\delta_1)).$$

The portfolio choice problem is

$$\begin{aligned}\max_{0\le a\le 1} V(p) &= \max_{0\le a\le 1}\max_{\delta_1\in\mathbb{R},\delta_2\in\mathbb{R}_+}\Big(\int b_\delta\circ u(f(s))\,dq(s)-\delta_1-\delta_2\eta\Big)\\ &= \max_{\delta_1\in\mathbb{R},\delta_2\in\mathbb{R}_+}\Big(\max_{0\le a\le 1}\Big(\int b_\delta\circ u(f(s)\,dq(s)\Big)-\delta_1-\delta_2\eta\Big).\end{aligned}\tag{30}$$

If for each $s\in S$ the distribution of $f(s)$ assigns probability one to a single point, then the solution to

$$\max_{0\le a\le 1}\int_S b_\delta\circ u(f(s))\,dq(s)$$

can be obtained by dynamic programming. The solution to the portfolio choice problem in (30) is completed by maximizing over $\delta_1\in\mathbb{R}$ and δ_2 in $\mathbb{R}_+$.

Maccheroni et al. (2006b) develop dynamic variational preferences. These are conditional preferences that impose consequentialism. The atemporal model is a special case of the dynamic model with one period of uncertainty—the atemporal model corresponds to preferences over one-step-ahead continuation plans. In the case of divergence preferences, the dynamic model can be applied to our two-period portfolio choice with state space $S=\{(h,h),(h,l),(l,h),(l,l)\}$, and the act f as in (11):

$$\begin{aligned}V_{r_0}(f) &= \min_{p\in\Delta(S)}\int_S u(f(r_0,r_1))\,dp(r_1\,|\,r_0)+\kappa D_\phi(p(\cdot\,|\,r_0)\|q(\cdot\,|\,r_0)) \qquad (r_0\in\{h,l\}),\\ V_0(f) &= \min_{p\in\Delta(S)}\int_S V_{r_0}(f)\,dp_0(r_0)+\kappa D_\phi(p_0\|q_0),\end{aligned}$$

where p is the joint distribution for (R_0,R_1) and p_0 is the marginal distribution for R_0. These dynamic preferences do not correspond to the ϕ-divergence preferences represented by (19) in the atemporal model, except when the ϕ-divergence is relative entropy, giving multiplier preferences. Maximizing the value functions gives

$$\begin{aligned}J_{r_0}(w) &= \max_{0\le x\le 1}\min_{p\in\Delta(S)}\int_S u(w[(r_1-r_f)x+r_f])\,dp(r_1\,|\,r_0)+\kappa D_\phi(p(\cdot\,|\,r_0)\|q(\cdot\,|\,r_0)),\\ J_0(w_0) &= \max_{0\le x\le 1}\min_{p\in\Delta(S)}\int_S J_{r_0}(w_0[(r_0-r_f)x+r_f])\,dp_0(r_0)+\kappa D_\phi(p_0\|q_0).\end{aligned}$$

In the case of multiplier preferences, this corresponds to the dynamic programming recursion in (13) for the atemporal model.

2.5. Smooth Ambiguity Preferences

Hansen & Miao note that the solution to the problem in (17) is a smooth ambiguity objective and a special case of Klibanoff et al. (2005) (KMM). The general form of smooth ambiguity (KMM) preferences values an act f as follows:

$$V(f)=\zeta^{-1}\Big[\int_\Theta \zeta\Big(\int_S u(f(s))\,d\pi_\theta(s)\Big)\,d\mu(\theta)\Big].\tag{31}$$

The state space $S=\Omega\times(0,1]$ and $f\colon S\to\mathcal{C}$ is a Savage act, where $\mathcal{C}\subset\mathbb{R}$ is a set of consequences. The set of Savage acts is denoted by $\mathcal{F}$. The space $(0,1]$ is introduced to model a rich set of lotteries as a set of Savage acts; there is an (objective) distribution on $(0,1]$ given by Lebesgue measure. An act $l\in\mathcal{F}$ is a lottery if l depends only on $(0,1]$. There is a preference ordering $\succsim$ over $\mathcal{F}$. π_θ is a prior distribution on S indexed by the parameter θ in the parameter space Θ. The function $\zeta\colon \mathbb{R} \to \mathbb{R}$ is strictly increasing. The distribution μ provides a prior on Θ. If the function ζ is linear, then the prior on priors reduces to a single prior $\int_\Theta \pi_\theta \, d\mu(\theta)$, but in general attitudes toward ambiguity are captured by a nonlinear ζ function.

KMM define second order acts $\mathfrak{f}\colon \Theta \to \mathcal{C}$ that associate an element of Θ to a consequence; $\mathfrak{F}$ denotes the set of second order acts and $\succsim^2$ is the decision maker's preference relation defined on $\mathfrak{F}$. Then KMM derive the representation in (31) from three assumptions. Assumption 1 is Expected Utility on Lotteries. This fixes a vNM utility function u, which is assumed to be strictly increasing. Assumption 2 is Subjective Expected Utility on Second Order Acts. This fixes a utility function ν, assumed to be strictly increasing, and a probability distribution μ on Θ such that for all $\mathfrak{f}, \mathfrak{g} \in \mathfrak{F}$,

$$\mathfrak{f} \succsim^2 \mathfrak{g} \quad \Longleftrightarrow \quad \int_\Theta \nu(\mathfrak{f}(\theta)) \, d\mu(\theta) \geq \int_\Theta \nu(\mathfrak{g}(\theta)) \, d\mu(\theta).$$

The set $\triangle(S)$ of probability distributions on S is indexed by the parameter $\theta \in \Theta$: $\triangle(S) = \{\pi_\theta : \theta \in \Theta\}$. An act f and a probability π_θ induce a probability distribution $\pi_{\theta,f}$ on consequences: with $B \subset \mathcal{C}$, $\pi_{\theta,f}(B) = \pi_\theta(f^{-1}(B))$. There is a lottery act with the distribution $\pi_{\theta,f}$ and certainty equivalent $c_{\theta,f}$, which is assumed to be the certainty equivalent of f given π_θ. Given $f \in \mathcal{F}$, $f^2 \in \mathfrak{F}$ denotes a second-order act associated with f: $f^2(\theta) = c_{\theta,f}$. Then Assumption 3 is Consistency with Preferences over Associated Second Order Acts: given $f, g \in \mathcal{F}$ and $f^2, g^2 \in \mathfrak{F}$,

$$f \succsim g \quad \Longleftrightarrow \quad f^2 \succsim^2 g^2.$$

These three assumptions imply that $\succsim$ is represented by (31) with $\zeta = \nu \circ u^{-1}$.

We can apply these preferences to our two-period portfolio choice problem. Let Ω be the state space in (10) with f equal to the act in (11) (so f depends only on Ω). Let the vNM utility function be $u(w) = w^{1-\gamma}/(1-\gamma)$. Let $\pi_\theta = \eta_\theta \times \lambda$, where η_θ is the distribution on Ω in (14) in which the returns (R_0, R_1) are i.i.d. conditional on θ, and λ is Lebesgue measure on $(0,1]$. For the strictly increasing function ζ use ζ_κ in (8). Let the parameter space Θ equal the unit interval $[0,1]$ with the prior distribution μ equal to a beta distribution with parameters α and β. Then the optimal portfolio weights can be obtained by maximizing the following objective function with respect to a_0, $a_1(h)$, $a_1(l)$ (which are part of the act f):

$$\begin{aligned} V(f) = -\kappa \log \int_0^1 \exp\Big(-\frac{1}{\kappa}\frac{1}{1-\gamma}\Big(\theta^2 f(h,h)^{1-\gamma} + \theta(1-\theta) f(h,l)^{1-\gamma} \\ + (1-\theta)\theta f(l,h)^{1-\gamma} + (1-\theta)^2 f(l,l)^{1-\gamma}\Big)\Big) \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha-1}(1-\theta)^{\beta-1} \, d\theta. \end{aligned} \tag{32}$$

$V(f)$ equals the objective function in (15), which we obtained using multiplier preferences with state space equal to the unit interval $[0,1]$, and we restricted the distribution for (R_0, R_1) to be i.i.d. conditional on θ.

An alternative appliction of KMM could continue to use Ω from (10) with f equal to the act in (11). But set π_θ equal to the predictive distribution in (1) with parameters α_θ and β_θ, with parameter space $\Theta = \{1, \ldots, K\}$ and a discrete prior distribution $\mu(k) = \xi_k$ (with $\xi_k \geq 0$ and $\sum_{k=1}^K \xi_k = 1$). These preferences do not correspond to multiplier preferences.

As $\kappa \to 0$, they exhibit maxmin expected utility behavior as in (4) and (5). See Klibanoff et al. (2005, Proposition 3).

Klibanoff et al. (2009) develop a model of recursive preferences that provides an intertemporal version of the smooth ambiguity model in KMM. In their approach, the atemporal model is a special case of the dynamic model with one period of uncertainty—the atemporal model corresponds to preferences over one-step-ahead continuation plans. The approach is similar to the way Maccheroni et al. (2006b) develop dynamic variational preferences from their atemporal model.

I have focused on modeling the decision maker's preferences over acts as a weak order $\succsim$ (complete and transitive). Some particular preference representations have been used: maxmin expected utility, multiplier preferences, divergence and constraint preferences, and smooth ambiguity preferences. These preference representations are related to Savage (1954) and Anscombe & Aumann (1963). Some other approaches to robust decisions are covered in the survey by Watson & Holmes (2016). The review by Stoye (2012) covers minimax regret. Manski (2004) and Stoye (2009) apply a minimax regret criterion in models of treatment choice.

3. PORTFOLIO CHOICE: EMPIRICAL WORK

In the preface to their monograph on *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Campbell & Viceira (2002) write (p. viii):

> There has always been a tension in economics between the attempt to describe the optimal choices of fully rational individuals ('positive economics') and the desire to use our models to improve people's imperfect choices ('normative economics'). The desire to improve the world with economics was well expressed by Keynes (1930): 'If economists could manage to get themselves thought of as humble, competent people, on a level with dentists, that would be splendid!' For much of the twentieth century, economists concentrated on improving economic outcomes through government economic policy; Keynes may have imagined the economist as orthodontist, intervening with the painful but effective tools of monetary and fiscal policy. Today dentists spend much of their time giving advice on oral hygiene; similarly, economists can try to provide useful advice to improve the myriad economic decisions that private individuals are asked to make. This book is an attempt at normative economics of this sort.

This section surveys some of the empirical work on normative portfolio choice. The framework is mainly subjective expected utility. Section 4 will consider how some of the theory from Section 2 on ambiguity could be applied in this context.

Barberis (2000) considers an investor with power utility over terminal wealth. There are two assets: Treasury bills and a stock index. The investor uses a vector autoregression (VAR) model that includes the excess return on the stock index and variables, such as the dividend-price ratio (dividend yield), that may be useful in predicting returns. It takes the form

$$z_t = a + Bx_{t-1} + \epsilon_t, \tag{33}$$

with $z_t' = (r_t, x_t')$, $x_t = (x_{1,t}, \ldots, x_{n,t})'$, and $\epsilon_t \sim$ i.i.d. $N(0, \Sigma)$. The first component of z_t is r_t, the continuously compounded excess stock return over month t (the rate of return on the stock portfolio minus the treasury bill rate, where both returns are continuously compounded). The remaining components of z_t are predictors of the stock return. For simplicity, the continuously compounded real monthly return on Treasury bills is treated as a constant r_f. Barberis considers both a buy and hold strategy and dynamic rebalancing. In both cases he examines the effect of parameter uncertainty. The empirical work uses postwar data on asset returns and the dividend yield, with 523 monthly observations from June 1952 through December 1995.

In the buy and hold case, the investor observes $\{z_t\}_{t=1}^{T}$ and chooses the allocation x to the stock index at $t = T$. If initial wealth $W_T = 1$, then end-of-horizon wealth is

$$W_{T+\hat{T}} = (1-x)\exp(r_f\hat{T}) + x\exp(r_f\hat{T} + r_{T+1} + \cdots + r_{T+\hat{T}}). \tag{34}$$

The investor's preferences over terminal wealth follow a power utility function u with coefficient of relative risk aversion equal to γ:

$$u(W) = \frac{W^{1-\gamma}}{1-\gamma}.$$

The investor's problem is

$$\max_x E_T\Big(\frac{[(1-x)\exp(r_f\hat{T}) + x\exp(r_f\hat{T} + R_{T+\hat{T}})]^{1-\gamma}}{1-\gamma}\Big),$$

where $R_{T+\hat{T}}$ denotes the cumulative excess return over $\hat{T}$ periods:

$$R_{T+\hat{T}} = r_{T+1} + r_{T+2} + \cdots + r_{T+\hat{T}}.$$

The investor calculates the expectation E_T conditional on his information set at time T.

In the version that ignores parameter uncertainty, the VAR parameters $\theta = (a, B, \Sigma)$ are estimated, and then the model is iterated forward with the parameters fixed at their estimated values. This generates a distribution for future stock returns conditional on a set of parameter values, which is denoted $p(R_{T+\hat{T}} \,|\, z, \hat{\theta})$, where $z = \{z_t\}_{t=1}^{T}$ is the data observed by the investor up until the start of his investment horizon. Then the investor's problem is

$$\max_x \int u(W_{T+\hat{T}})p(R_{T+\hat{T}} \,|\, z, \hat{\theta})\, dR_{T+\hat{T}}. \tag{35}$$

In order to allow for parameter uncertainty, a single prior distribution for θ is specified. This prior is intended to be relatively uninformative, so that it can be dominated by sample evidence. The likelihood based on (33) and the prior imply a posterior distribution $p(\theta \,|\, z)$ and a predictive distribution for long-horizon returns:

$$p(R_{T+\hat{T}} \,|\, z) = \int p(R_{T+\hat{T}} \,|\, z, \theta)p(\theta \,|\, z)\, d\theta.$$

Then the investor's problem is

$$\max_x \int u(W_{T+\hat{T}})p(R_{T+\hat{T}} \,|\, z)\, dR_{T+\hat{T}}. \tag{36}$$

Without using predictor variables (so $z_t = r_t$ and x_t is null) and ignoring parameter uncertainty (as in (35)), the optimal portfolio weight on the risky asset is approximately independent of the investment horizon. Barberis notes the similarity to the result in Samuelson (1969), which shows that with power utility and independent and identically distributed (i.i.d.) returns, the optimal allocation is independent of the horizon. This, however, is for an investor who optimally rebalances his portfolio, rather than the buy and hold investor considered here. Allowing for parameter uncertainty (as in (36)), the allocation to stocks falls as the horizon increases. The magnitude of this effect is substantial. For an investor using the full data set and with a coefficient of relative risk aversion equal to 5, the difference in allocation at a ten-year horizon compared with a one-year horizon is roughly 10 percentage points. If the investor only uses data from 1986 to 1995, the difference is 35 percentage points.

Now consider including the dividend yield as a predictor variable x_t in the VAR. Ignoring parameter uncertainty, the optimal allocation to stocks for a long-horizon investor is much higher than for a short-horizon investor. When the uncertainty about the parameters is taken into account, the long-horizon allocation is again higher than the short-horizon allocation; but the difference between the long and short horizons is not nearly as great as when estimation risk is ignored.

Barberis also considers dynamic allocation in which the investor optimally rebalances over his investment horizon. Consider the case without predictor variables, so $z_t = r_t = a + \epsilon_t$ in (33) is i.i.d. conditional on the parameters a and Σ. Now allowing for parameter uncertainty involves learning because the uncertainty about the parameters changes over time. As new data are observed, the investor updates his posterior distribution for the parameters. The investor anticipates this learning, and it affects his portfolio holdings. This corresponds to the two-period problem with a single prior that we discussed in Section 2.1. The investor's problem corresponds to (2), which can be solved by dynamic programming as in (3). Barberis uses dynamic programming to calculate the optimal allocation to stocks at T for horizons $\hat{T}$ varying from 1 to 10 years. The result is that the investor who acknowledges the parameter uncertainty allocates less to stocks at longer horizons. The magnitude of the effect is substantial and similar to the results for the buy and hold strategy.

Xia (2001) works with a continuous-time model based on Brownian motion. This leads to closed-form formulas for optimal portfolios with learning in some special cases. Maenhout (2004, 2006) works with a continuous-time model based on Brownian motion. Using results from Anderson et al. (2003), he obtains closed-form solutions for robust portfolio rules in some special cases.

Kandel and Stambaugh (1996) consider a problem similar to that of Barberis (2000), but with a one month horizon ($\hat{T} = 1$) and a potentially large number n of predictor variables. They are interested in providing a metric to assess the economic significance of the regression evidence on stock-return predictability. They use the perspective of a single-prior Bayesian investor who uses the sample evidence (with a likelihood function based on (33)) to update prior beliefs about the regression parameters. The investor then uses these revised beliefs to compute the optimal asset allocation. They specify a prior that is intended to be relatively uninformative, and also an informative prior that is weighted against return predictability. They find that the economic significance of the sample evidence need not correspond to standard statistical measures. An investor can assign an important role to the predictor variables even though the regression results produce a large p-value for the null hypothesis that the coefficients on the predictor variables are jointly equal to zero. This is particularly relevant with a large number, say 25, of predictor variables. The investor's allocation decision does not involve accepting or rejecting a specific hypothesis. The investor's problem is to select a portfolio, not a hypothesis.

Stambaugh (1999) also considers a problem similar to that of Barberis (2000), with a focus on a single predictor variable equal to the dividend yield on the aggregate stock market portfolio. The dividend yield is highly persistent with estimated autoregression coefficient close to one. This leads to sharp contrasts between frequentist and Bayesian inference. Stambaugh explores this, providing empirical counterparts to issues raised by Sims (1988) and Sims & Uhlig (1991). Stambaugh develops predictive distributions, which incorporate "estimation risk" arising from parameter uncertainty. These are used to calculate an optimal portfolio for a buy-and-hold investor facing a stocks-versus-cash allocation decision. He considers investment horizons ranging from one month to 20 years. He examines sensitivity to conditioning on the initial observation in forming the likelihood function (based on (33)) versus treating the initial observation as a draw from the ergodic distribution. He also examines sensitivity to alternative prior specifications that are intended to be uninformative.

Pástor & Stambaugh (2012) base their likelihood function on the following model:

$$\begin{aligned} r_{t+1} &= \mu_t + u_{t+1}, \\ x_{t+1} &= \theta + A x_t + v_{t+1}, \\ \mu_{t+1} &= (1-\beta) + \beta\mu_t + w_{t+1}. \end{aligned} \tag{37}$$

Their annual data consists of observations for the 206-year period from 1802 through 2007, as compiled by Siegel (1992, 2008). The return r_t is the annual real log return on the U.S. equity market, and x_t contains three predictors: the dividend yield on U.S. equity, the first difference in the long-term high-grade bond yield, and the difference between the long-term bond yield and the short-term interest rate. The variable μ_t is not observed. It is motivated by considering the possibility of an information set $\mathcal{F}_t$ that includes the observed data $\{r_t, x_t\}_{t=1}^T$ and additional predictor variables that are not observed by the investor. Then $\mu_t = E(r_{t+1} \,|\, \mathcal{F}_t)$. Conditional on $\mathcal{F}_t$, the innovation vector $(u_{t+1}, v_{t+1}, w_{t+1})$ is i.i.d. $N(0, \Sigma)$. The observed predictor variables x_t are related to the latent μ_t because v_t and w_t are correlated. In this model the mean of r_{t+1} conditional on the observed $\{r_s, x_s\}_{s=1}^t$ depends in a parsimonious way on the entire history, not just on (r_t, x_t).

The authors specify a range of informative prior distributions. A key prior distribution is the one on the correlation ρ_{uw} between u_t and w_t. Their benchmark prior has 97% of its mass below zero. This prior is based on the argument in Pástor and Stambaugh (2009) that the correlation between innovations in return and expected return are likely to be negative. The authors use the likelihood function based on their model in (37) and their prior distributions to calculate optimal stock allocations for an investor in a target-date fund. The investor's horizon is K years, and his utility for end-of-horizon wealth W_K is $W_K^{1-\gamma}/(1-\gamma)$. The investor commits to a predetermined investment strategy in which the stock allocation evolves linearly from the first-period allocation x_1 to the final-period allocation x_K. When parameter uncertainty is ignored, the parameters in (37) are treated as known and equal to their posterior means. In that case the initial allocation increases steadily as the investment horizon lengthens, increasing from 30% at the one-year horizon to about 85% at long horizons of 25 or 30 years (with $\gamma = 8$). The results are quite different when the predictive distribution is used to incorporate parameter uncertainty. The initial allocation increases from 30% at the one year horizon to 57% at the 30-year horizon.

Campbell & Viceira (2002) provide insights into how an individual investor would best allocate wealth into broad asset classes over a lifetime. The authors use approximate analytical solutions to long-term portfolio choice problems. This provides analytical insights in models that fall outside the limited class that can be solved exactly. One of their models allows consumption at every date. The intertemporal budget constraint is that wealth next period equals the portfolio return times reinvested wealth, that is, wealth today less what is subtracted for consumption:

$$W_{t+1} = (1 + R_{p,t+1})(W_t - C_t).$$

Preferences over random consumption streams are defined recursively:

$$U_t = \left[(1-\delta)C_t^{1-\rho} + \delta D_t^{1-\rho}\right]^{\frac{1}{1-\rho}}, \quad D_t = (E_t U_{t+1}^{1-\gamma})^{\frac{1}{1-\gamma}},$$

where δ equals the time discount factor, $\psi = 1/\rho$ equals the intertemporal elasticity of substitution, and γ is related to risk aversion. These preferences were developed by Epstein & Zin (1989, 1991) and Weil (1989) using the theoretical framework of Kreps & Porteus (1978).

In Chapter 4, *Is the Stock Market Safer for Long-Term Investors?*, the authors use their general framework to investigate how investors should allocate their portfolios among three assets: stocks, nominal bonds, and nominal treasury bills. Investment opportunities are described using a VAR system that includes short-term ex post real interest rates, excess stock returns, excess bond returns, and variables that have been identified as return predictors by empirical research: the short-term nominal interest rate, the dividend-price ratio, and the yield spread between long-term bonds and Treasury bills. The annual data covers 1890 to 1998. Its source is the data used in Grossman & Shiller (1981), updated following the procedures of Campbell (1999). Point estimates from the VAR are used in the analytic formulas for optimal portfolios. A range of different values for γ are used, assuming $\psi = 1$ and $\delta = .92$ in annual terms. The investor optimally rebalances the portfolio each period. The solutions do not impose constraints that might prevent short selling or borrowing to invest in risky assets. At $\gamma = 5$, the stock allocation is 67%, the bond allocation is 91%, and the cash position is -58% to finance stock and bond positions that exceed 100% of the portfolio. As risk aversion increases above 5, the demand for bonds increases and accounts for almost the entire portfolio of extremely conservative long-term investors.

Campbell & Viceira develop asset allocation models with labor income in Chapter 6, and they consider the role of labor income risk and precautionary savings. Empirical results from the Panel Study of Income Dynamics (PSID) are used to calibrate the models. In Chapter 7, they develop a life-cycle model of consumption and portfolio choice. They use the PSID to measure differences in the stochastic structure of the labor income process across industries, and differences between self-employed and non-self-employed households. They examine the effects of these differences, and of other sources of investor heterogeneity, on optimal consumption and portfolio choice. These models and questions suggest a rich set of decision problems in which a variety of data sets can be used, including administrative data. We could also include human capital investment decsions that individuals make for themselves and their children, location decisions, and health decisions involving treatment choice.

4. WORKING WITH DIVERGENCE AND CONSTRAINT PREFERENCES

This section discusses how the analysis in Barberis (2000) could work with divergence and constraint preferences. We shall also consider including consumption choices as in Campbell and Viceira (2002). First consider a buy and hold strategy. The reference model q is based on the vector autoregression (VAR) in (33) and a prior distribution π for the parameters $\theta = (a, B, \Sigma)$ in the parameter space Θ. Let $z = (z^{(1)}, z^{(2)})$ with $z^{(1)} = \{z_t\}_{t=1}^{T}$ and $z^{(2)} = \{z_t\}_{t=T+1}^{T+\hat{T}}$. The VAR provides a conditional density $\lambda(z \,|\, \theta)$ for $z \in \mathcal{Z}$. Set the state space $S = \mathcal{Z} = \mathcal{Z}^{(1)} \times \mathcal{Z}^{(2)}$. Then for $A \subset \mathcal{Z}$,

$$q(A) = \int_A \int_\Theta \lambda(z \,|\, \theta)\, dz\, d\pi(\theta).$$

The investor observes $z^{(1)}$ and chooses the allocation x to the stock index at $t = T$ as a function of $z^{(1)}$: $x = a(z^{(1)}) \in [0, 1]$. The corresponding act is based on end-of-horizon wealth as in (34) (with $W_T = 1$):

$$f(z) = f(a(z^{(1)}), z^{(2)}) = (1 - a(z^{(1)})) \exp(r_f \hat{T}) + a(z^{(1)}) \exp(r_f \hat{T} + r_{T+1} + \cdots + r_{T+\hat{T}}). \quad (38)$$

The investor's problem is

$$\max_{0 \le a(\cdot) \le 1} \min_{p \in \Delta(\mathcal{Z})} \int_{\mathcal{Z}} u(f(z))\, dp(z) + \kappa D_\phi(p \,\|\, q). \quad (39)$$

From Section 2.4,

$$\min_{p \in \Delta(\mathcal{Z})} \int_{\mathcal{Z}} u(f(z))\, dp(z) + \kappa D_\phi(p \,\|\, q) = \max_{\delta \in \mathbb{R}} -\kappa \int_{\mathcal{Z}} \phi^*(-\frac{u(f(z)) + \delta}{\kappa})\, dq(z) - \delta.$$

For a given value of δ, define the increasing, concave function $b_\delta : \mathbb{R} \to [-\infty, \infty)$ by

$$b_\delta(x) = -\kappa \phi^*(-\frac{x + \delta}{\kappa}).$$

Then (39) is equivalent to

$$\max_{\delta \in \mathbb{R}} \Big(\max_{0 \le a(\cdot) \le 1} \Big(\int_{\mathcal{Z}} b_\delta \circ u(f(z))\, dq(z) \Big) - \delta \Big). \quad (40)$$

For each $z \in \mathcal{Z}$, the distribution of $f(z)$ puts probability one on a single point. So we can replace the utility function u by $b_\delta \circ u$ and apply SEU preferences using the reference distribution q:

$$\max_{0 \le a(\cdot) \le 1} \int_{\mathcal{Z}} b_\delta \circ u(f(z))\, dq(z) = \int_{Z^{(1)}} \Big(\max_{x \in [0,1]} \int_{Z^{(2)}} b_\delta \circ u(f(x, z^{(2)}))\, dq(z^{(2)} \,|\, z^{(1)}) \Big)\, dq(z^{(1)}).$$

Then the solution to (40) is completed by maximizing over δ in $\mathbb{R}$. In the case of multiplier preferences, replacing $D_\phi(p \,\|\, q)$ by $R(p \,\|\, q)$, the solution a^* to (39) simplifies to

$$a^*(z^{(1)}) = \arg \max_{x \in [0,1]} \int_{Z^{(2)}} \zeta_\kappa \circ u(f(x, z^{(2)}))\, dq(z^{(2)} \,|\, z^{(1)}).$$

As in Section 2.4, we can use the solutions for a range of values of the multiplier κ to provide solutions using the associated constraint preferences for a range of values for the constraint parameter η in (27). Then the constraint parameter η will fix the convex set $\mathcal{C}$ of distributions for MEU preferences. The choice of η can be based on considering the set of implied predictive distributions for portfolio returns. Following Good (1952), only reasonable predictive distributions should be included.

Now consider applying robustness to the prior $\pi(\theta)$, maintaining the conditional density $\lambda(z \mid \theta)$. Set the state space S equal to the parameter space Θ. Using the notation for $f(z)$ in (38), the act $g(\theta)$ is an objective distribution given by (for $A \subset \mathbb{R}$)

$$g(\theta)(A) = \int_{f^{-1}(A)} \lambda(z \mid \theta)\, dz.$$

Evaluating the vNM utility function u at $g(\theta)$ gives

$$u(g(\theta)) = \int_{\mathcal{Z}} u(f(z))\lambda(z \mid \theta)\, dz.$$

With multiplier preferences:

$$\min_{p \in \Delta(\Theta)} \int_{\Theta} u(g(\theta))\, dp(\theta) + \kappa R(p \,\|\, \pi) = \zeta_\kappa^{-1}\Big(\int_{\Theta} \zeta_\kappa(u(g(\theta)))\, d\pi(\theta)\Big).$$

Now it matters that $\zeta_\kappa \circ u$ is not a vNM utility function. The optimal portfolio rule solves

$$\max_{0 \le a(\cdot) \le 1} \int_{\Theta} \zeta_\kappa\Big(\int_{\mathcal{Z}} u(f(a(z^{(1)}), z^{(2)}))\lambda(z \mid \theta)\, dz\Big)\, d\pi(\theta).$$

This does not correspond to applying SEU preferences with a modified utility function.

Consider allowing for consumption and portfolio choice at dates $T, T+1, ..., T+\hat{T}-1$, as in Campbell & Viceira (2002). Let $z^t = \{z_\tau\}_{\tau=1}^{t}$. At date T, wealth w_T is given and consumption is a function of w_T and the history z^T: $c_T = c_T(w_T, z^T)$. Likewise, the stock allocation is $a_T = a_T(w_T, z^T)$. Then wealth is defined recursively:

$$w_{t+1} = [w_t - c_t(w_t, z^t)][(1 - a_t(w_t, z^t))\exp(r_f) + a_t(w_t, z^t)\exp(r_f + r_{t+1})]$$

for $t = T, T+1, \ldots, T+\hat{T}-1$. Define $T' = T + \hat{T}$. Multiplier preferences are given by

$$U = \min_{p \in \Delta(\mathcal{Z})} \int \Big(\sum_{t=T}^{T'-1} u_t(c_t) + u_{T'}(w_{T'})\Big)\, dp + \kappa R(p \,\|\, q).$$

We shall see that these preferences are recursive. First we need to decompose $R(p \,\|\, q)$ into conditional and marginal components.

Let $(Z_1, \ldots, Z_{T'})$ denote the random vector with distribution p, and let p^t denote the marginal distribution of $(Z_1, \ldots, Z_t)$. Let $p^+(\cdot \mid z^t)$ denote the conditional distribution of Z_{t+1} given $(Z_1, \ldots, Z_t) = z^t$. Let $h_t = dp_t/dq_t$ and apply iterated expectations:

$$\begin{aligned}
R(p_t \,\|\, q_t) &= \int \log[h_t(z^t)]\, dp^t \\
&= \int \log[h_t(z_t \mid z^{t-1})h_{t-1}(z^{t-1})]\, dp^t \\
&= \int \Big(\int \log[h_t(z_t \mid z^{t-1})]\, dp^+(z_t \mid z^{t-1})\Big)\, dp^{t-1} + \int \log[h_{t-1}(z^{t-1})]\, dp^{t-1} \\
&= \int R(p^+(\cdot \mid z^{t-1}) \,\|\, q^+(\cdot \mid z^{t-1}))\, dp^{t-1} + R(p^{t-1} \,\|\, q^{t-1}).
\end{aligned}$$

Note that

$$\begin{aligned}
&\int \Bigg(\sum_{t=T}^{T'-1} u_t(c_t) + \Bigg(\int u_{T'}(w_{T'})\, dp^+(z_{T'} \,|\, z^{T'-1}) \\
&\quad + \kappa R(p^+(\,\cdot\,|\, z^{T'-1}) \,\|\, q^+(\,\cdot\,|\, z^{T'-1})) \Bigg) \Bigg) \, dp^{T'-1} + \kappa R(p^{T'-1} \,\|\, q^{T'-1}) \\
&\geq \int \Bigg(\sum_{t=T}^{T'-1} u_t(c_t) + \min_{p^+(\,\cdot\,|\, z^{T'-1})} \Bigg(\int u_{T'}(w_{T'})\, dp^+(z_{T'} \,|\, z^{T'-1}) \\
&\quad + \kappa R(p^+(\,\cdot\,|\, z^{T'-1}) \,\|\, q^+(\,\cdot\,|z^{T'-1})) \Bigg) \Bigg) \, dp^{T'-1} + \kappa R(p^{T'-1} \,\|\, q^{T'-1}).
\end{aligned}$$

Suppose that $\hat{T} \geq 2$ and define the value function

$$V_{T'-1} = u_{T'-1}(c_{T'-1}) + \zeta_\kappa^{-1} \left(\int \zeta_\kappa(u_{T'}(w_{T'}))\, dq^+(z_{T'} \,|\, z^{T'-1}) \right).$$

Then

$$U = \min_{p^{T'-1}} \int \left(\sum_{t=T}^{T'-2} u_t(c_t) + V_{T'-1} \right) dp^{T'-1} + \kappa R(p^{T'-1} \,\|\, q^{T'-1}).$$

Recursively define the value functions

$$V_t = u_t(c_t) + \zeta_\kappa^{-1} \left(\int \zeta_\kappa(V_{t+1})\, dq^+(z_{t+1} \,|\, z^t) \right) \qquad (t = T'-2, T'-3, \ldots, T). \tag{41}$$

Make the inductive assumption: for some $t \geq T$,

$$U = \min_{p^{t+1}} \int \left(\sum_{\tau=T}^{t} u_\tau(c_\tau) + V_{t+1} \right) dp^{t+1} + \kappa R(p^{t+1} \,\|\, q^{t+1}).$$

Note that

$$\begin{aligned}
&\int \Bigg(\sum_{\tau=T}^{t} u_\tau(c_\tau) + \Bigg(\int V_{t+1}\, dp^+(z_{t+1} \,|\, z^t) + \kappa R(p^+(\,\cdot\,|\, z^t) \,\|\, q^+(\,\cdot\,|\, z^t)) \Bigg) \Bigg) \, dp^t \\
&\quad + \kappa R(p^t \,\|\, q^t) \\
&\quad \geq \int \Bigg(\sum_{\tau=T}^{t} u_\tau(c_\tau) + \min_{p^+(\,\cdot\,|\, z^t)} \Bigg(\int V_{t+1}\, dp^+(z_{t+1} \,|\, z^t) + \kappa R(p^+(\,\cdot\,|\, z^t) \,\|\, q^+(\,\cdot\,|\, z^t)) \Bigg) \Bigg) \, dp_t \\
&\quad + \kappa R(p^t \,\|\, q^t).
\end{aligned}$$

Then

$$\begin{aligned}
U &= \min_{p^t} \int \left(\sum_{\tau=T}^{t-1} u_\tau(c_\tau) + V_t \right) dp^t + \kappa R(p^t \,\|\, q^t) \qquad (t = T'-1, T'-2, \ldots, T+1) \\
&= \min_{p^T} \int V_T \, dp^T + \kappa R(p^T \,\|\, q^T).
\end{aligned}$$

The value function V_t depends upon the policy functions $\{c_\tau\}_{\tau=t}^{T'-1}$, $\{a_\tau\}_{\tau=t}^{T'-1}$, and on wealth w and the history z^t: $V_t = V_t(w, z^t; \{c_\tau\}_{\tau=t}^{T'-1}, \{a_\tau\}_{\tau=t}^{T'-1})$. These recursive preferences have the important property of probabilistic sophistication with respect to the reference distribution q: if acts f and g imply the same distribution of outcomes under q, then $f \sim g$; see Macheroni et al. (2006a) and Machina & Schmeidler (1992).

Given the recursion for the value function in (41), we can apply dynamic programming to obtain optimal policy functions. Starting at the end of the horizon, construct the maximized value function at $T' - 1$:

$$J_{T'-1}(w, z^{T'-1}) = \max_{0 \le x_1 \le w, 0 \le x_2 \le 1} \left[u_{T'-1}(x_1) + \zeta_\kappa^{-1} \left(\int \zeta_\kappa(u_{T'}(\tilde{w}))\, dq^+(z_{T'} \mid z^{T'-1}) \right) \right]$$

with

$$\tilde{w} = (w - x_1)[(1 - x_2)\exp(r_f) + x_2 \exp(r_f + r_{T'})].$$

Then, working backwards, recursively construct the maximized value functions for the other dates:

$$J_t(w, z^t) = \max_{0 \le x_1 \le w, 0 \le x_2 \le 1} \left[u_t(x_1) + \zeta_\kappa^{-1} \left(\int \zeta_\kappa(J_{t+1}(\tilde{w}, z^{t+1}))\, dq^+(z_{t+1} \mid z^t) \right) \right] \tag{42}$$

$(t = T' - 2, T' - 3, \ldots, T)$ with

$$\tilde{w} = (w - x_1)[(1 - x_2)\exp(r_f) + x_2 \exp(r_f + r_{t+1})].$$

The optimal policy functions c_t^* and a_t^* evaluated at (w, z^t) are obtained from the solutions (arg max) for x_1 and x_2 in (42).

With dynamic divergence preferences, the atemporal model is applied to the one-step ahead continuation plans to form conditional preferences. The value function at $T' - 1$ is

$$V_{T'-1} = u_{T'-1}(c_{T'-1}) + \max_{\delta \in \mathbb{R}} -\kappa \left(\int \phi^*(-\frac{u_{T'}(w_{T'}) + \delta}{\kappa})\, dq^+(z_{T'} \mid z^{T'-1}) \right) - \delta,$$

where $\phi^* : \mathbb{R} \to (-\infty, \infty]$ is the convex conjugate of ϕ:

$$\phi^*(x) = \sup_{y \ge 0}[xy - \phi(y)].$$

The value function recursion in (41) becomes

$$V_t = u_t(c_t) + \max_{\delta \in \mathbb{R}} -\kappa \left(\int \phi^*(-\frac{V_{t+1} + \delta}{\kappa})\, dq^+(z_{t+1} \mid z^t) \right) - \delta. \tag{43}$$

As in the multplier case, these recursive preferences have probabilistic sophistication with respect to the reference distribution q. Given the recursion for the value function in (43), we can apply dynamic programming to obtain optimal policy functions.

Strzalecki (2013) shows how these recursive preferences can be reinterpreted. Instead of ambiguity aversion, they can reflect a preference for the timing of the resolution of uncertainty, as in Kreps and Porteus (1978).

5. CRITIQUE

In his article on "Pitfalls to a Minimax Approach to Model Uncertainty," Sims (2001) wrote (p. 51):

> In fact, because it violates the sure-thing principle, most people, and certainly most policymakers, would be likely to alter behavior fitting the maximin theory if they were shown certain consequences of it. For example, because maximin expected utility behavior, if applied de novo to each of a sequence of choice sets, can imply behavior consistent with no single set of probabilistic prior beliefs, it can allow a Dutch Book, a situation where someone agrees to a set of bets that causes him to lose money with probability 1.

Multiplier preferences satisfy the sure-thing principle on a given state space, but this does not necessarily address Sims' critique. In Section 2.3 we imposed the restriction that the alternative distributions p for (R_0, R_1) be independent and identically distributed conditional on θ, and we used $\{\theta : 0 \leq \theta \leq 1\}$ as the state space. In that case, the multiplier preferences were not recursive with respect to the filtration generated by (R_0, R_1).

A similar issue arose with the Ellsberg three-color problem when we restricted the alternative distributions p to assign probability 1/3 to drawing the A ball. The state space was $\{\theta : 0 \leq \theta \leq 1\}$. Then our solution exhibited typical Ellsberg behavior, which violates the sure-thing principle applied to the state space $\{A, B, C\}$ in Table 1 and Figure 3.

On the other hand, if we do not restrict the predictive distribution for (R_0, R_1), then we obtain recursive preferences and can solve for the optimal portfolio using dynamic programming. Likewise, in Section 4, when we do not restrict the predictive distribution for $\{Z_t\}_{t=1}^{T+\hat{T}}$, the multiplier and divergence preferences are recursive with respect to the filtration generated by $\{Z_t\}_{t=1}^{T+\hat{T}}$. These preferences satisfy consequentialism, which is related to the sure-thing principle.

The lack of a single set of probabilistic beliefs can arise if one uses the least-favorable prior as the beliefs. This least-favorable prior depends on the feasible set of actions and will change if the problem changes. An alternative is to use the reference distribution q to provide a single set of probabilistic beliefs. The multiplier and divergence preferences are probabilistically sophisticated with respect to q for a given state space. This does not, however, necessarily address Sims' critique. Consider the Ellsberg three-color problem, where we restrict the alternative distributions p to assign probability 1/3 to drawing the A ball, and we use $\{\theta : 0 \leq \theta \leq 1\}$ as the state space. The multiplier preferences in Section 2.3 are probabilistically sophisticated with respect to this state space, but they exhibit typical Ellsberg behavior. So they are not probabilistically sophisticated with respect to the state space $\{A, B, C\}$ (Machina & Schmeidler (1992, p. 752)).

6. CONCLUSION

My preference is to stay within the SEU framework. Then multiplier preferences can be used in a sensitivity analysis, with the reference distribution q based on SEU. The multiplier parameter κ can be chosen to ensure that the least-favorable distribution is plausible. The least-favorable prior can focus attention on parts of the reference distribution that are particularly important for the decision problem. The sensitivity analysis can provide protection against specifying a reference distribution that has unintended, dogmatic aspects.

If the reference distribution is built up from a likelihood function $\lambda(z \,|\, \theta)$ and a prior distribution π on the parameter space Θ, then one possibility is to hold the likelihood function fixed and apply the sensitivity analysis only to the prior distribution π. We saw in Section 5 that multiplier preferences in this case are particularly vulnerable to a normative critique. We are on firmer ground when the alternative distributions p do not restrict the predictive distribution, so that the likelihood function λ and the prior π on the parameter space are both subject to the sensitivity analysis. Then the multiplier preferences are recursive and satisfy consequentialism with respect to the (natural) filtration generated by the observations $\{Z_t\}_{t=1}^{T+\hat{T}}$.

Divergence and dynamic divergence preferences can also be used in a sensitivity analysis.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

For comments and discussions, I am grateful to Isaiah Andrews, Timothy Christensen, Max Kasy, Thomas Knox, Mikkel Plagborg-Møller, Ashesh Rambachan, Brit Sharoni, Tomasz Strzalecki, and Elie Tamer.

LITERATURE CITED

Anderson EW, Hansen LP, Sargent TJ. 2003. A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *J. Eu. Econ. Assoc.* 1:68–123.

Anscombe FJ, Aumann RJ. 1963. A definition of subjective probability. *Ann. Math. Stat.* 34:199–205

Barberis N. 2000. Investing for the long run when returns are predictable. *J. Financ.* 55:225–264

Blackwell D, Girshick MA. 1954. *Theory of Games and Statistical Decisions.* NY: Wiley

Campbell JY. 1999. Asset prices, consumption, and the business cycle. In *Handbook of Macroeconomics*, Vol. 1, ed. JB Taylor, M. Woodford, pp. 1231–303. Amsterdam: North Holland

Campbell JY, Viceira LM. 2002. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors.* Oxford: Oxford University Press

Chamberlain G. 2000a. Econometrics and decision theory. *J. Econom.* 95:255–83

Chamberlain G. 2000b. Econometric applications of maxmin expected utility. *J. Applied Econom.* 15:625–44

Christensen T, Connault B. 2019. Counterfactual sensitivity and robustness. arXiv:1904.00989v2 [econ.EM]

Diaconis P, Skyrms B. 2018. *Ten Great Ideas About Chance.* Princeton: Princeton University Press

Dupuis P, Ellis RS. 1997. *A Weak Convergence Approach to the Theory of Large Deviations.* NY: Wiley

Ellsberg D. 1961. Risk, ambiguity, and the savage axioms. *Q. J. Econ.* 75:643–69

Epstein LG, Zin SE. 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. *Econometrica.* 57:937–69

Epstein LG, Zin SE. 1991. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: an empirical analysis. *J. Polit. Econ.* 99:263–86

Epstein LG, Schneider M. 2003. Recursive multiple-priors. *J. Econ. Theory.* 113:1–31

Ferguson TS. 1967. *Mathematical Statistics: A Decision Theoretic Approach.* NY: Academic Press

Gilboa I, Schmeidler D. 1989. Maxmin expected utility with non-unique prior. *J. Math. Econ.* 18:141–53.

Good IJ. 1952. Rational decisions. *J. Royal Stat. Soc. Ser. B.* 14:107–14

Grossman SJ, Shiller RJ. 1981. The determinants of the variability of stock market prices. *Am. Econ. Rev.* 71:222–7

Gumen A, Savochkin A. 2013. Dynamically stable preferences. *J. Econ. Theory.* 148:1487–1508

Hanany E, Klibanoff P. 2007. Updating preferences with multiple priors. *Theor. Econ.* 2:261–98

Hansen LP, Sargent TJ. 2001. Robust control and model uncertainty. *Am. Econ. Rev.* 91:60–6

Hansen LP, Sargent TJ, Turmuhambetova G, Williams N. 2006. Robust control and model misspecification. *J. Econ. Theory.* 128:45–90

Hansen LP, Sargent TJ. 2008. *Robustness.* Princeton: Princeton University Press

Hansen LP, Sargent TJ. 2015. *Uncertainty Within Economic Models.* Singapore: World Scientific

Hansen LP, Miao J. 2018. Aversion to ambiguity and model misspecification in dynamic stochastic environments. *Proc. Natl. Acad. Sci. USA* 115:9163–68

Kandel S, Stambaugh RF. 1996. On the predictability of stock returns: an asset-allocation perspective. *J. Financ.* 51:385–424

Keynes JM. 1930. Economic possibilities for our grandchildren. In JM Keynes, 1932, *Essays in Persuasion.* NY: Harcourt Brace

Kitamura Y, Otsu T, Evdokimov K. 2013. Robustness, infinitesimal neighborhoods, and moment conditions. *Econometrica.* 81:1185–201

Klibanoff P, Marinacci M, Mukerji S. 2005. A smooth model of decision making under uncertainty. *Econometrica.* 73:1849–92

Klibanoff P, Marinacci M, Mukerji S. 2009. Recursive smooth ambiguity preferences. *J. Econ. Theory.* 144:930–76

Knox TA. 2003a. *Learning How to Invest when Returns are Uncertain.* PhD Thesis, Harvard University

Knox TA. 2003b. Foundations for learning how to invest when returns are uncertain. Unpublished manuscript, Univ. of Chicago

Kreps DM, Porteus EL. 1978. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica.* 46:185–200

Kreps DM, Porteus EL. 1979 Dynamic Choice Theory and Dynamic Programming. *Econometrica.* 47:91–100

Luenberger DG. 1969. *Optimization by Vector Space Methods.* NY: Wiley

Maccheroni F, Marinaci M, Rustichini A. 2006a. Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica.* 74:1447–98

Maccheroni F, Marinacci M, Rustichini A. 2006b. Dynamic Variational Preferences. *J. Econ. Theory.* 128:4–44

Machina MJ. 1989. Dynamic consistency and non-expected utility models of choice under uncertainty. *J. Econ. Lit.* 27:1622–68

Machina MJ, Schmeidler D. 1992. A more robust definition of subjective probability. *Econometrica.* 60:745–80

Maenhout PJ. 2004. Robust portfolio rules and asset pricing. *Rev. Financial Stud.* 17:951–83

Maenhout PJ. 2006. Robust portfolio rules and detection-error probabilities for a mean-reverting risk premium. *J. Econ. Theory.* 128:136–63

Manski CF. 2004. Statistical treatment rules for heterogeneous populations. *Econometrica.* 72:1221–46

Newey WK, Smith RJ. 2004. Higher order properties of GMM and generalized empirical likelihood estimators. *Econometrica.* 72:219–55

Pástor L, Stambaugh RF. 2009. Predictive systems: living with imperfect predictors. *J. Financ.* 64:1583–628

Pástor L, Stambaugh RF. 2012. Are stocks really less volatile in the long run? *J. Financ.* 67:431–77

Ramsey FP. 1926. Truth and probability. In FP Ramsey, 1931, *The Foundations of Mathematics and other Logical Essays*, Chap. 7, ed. RB Braithwaite, pp. 156–98. London/NY: Routledge

Samuelson PA. 1969. Lifetime portfolio selection by dynamic stochastic programming. *Rev. Econ. Stat.* 51:239–46

Savage LJ. 1954. *The Foundations of Statistics.* NY: Wiley

Siegel JJ. 1992. The real rate of interest from 1800–1990: a study of the U.S. and the U.K. *J.*

Monetary Econ. 29:227–52
Siegel JJ. 2008. *Stocks for the Long Run.* NY: McGraw Hill
Sims CA. 1988. Bayesian skepticism on unit root econometrics. *J. Econ. Dyn. Control.* 12:463–74
Sims CA, Uhlig H. 1991. Understanding unit rooters: a helicopter tour. *Econometrica.* 59:1591–99
Stambaugh RF. 1999. Predictive regressions. *J. Financial Econ.* 54:375–421
Stoye J. 2009. Minimax regret treatment choice with finite samples. *J. Econom.* 151:70–81
Stoye J. 2012. New perspectives on statistical decisions under ambiguity. *Annu. Rev. Econ.* 4:257–82
Strzalecki, T. 2011. Axiomatic foundations for multiplier preferences. *Econometrica.* 79:47–73
Strzalecki, T. 2013. Temporal resolution of uncertainty and recursive models of ambiguity aversion. *Econometrica.* 81:1039–74
Watson J, Holmes C. 2016. Approximate models and robust decisions. *Stat. Sci.* 31:465–89
Weil P. 1989. The equity premium puzzle and the risk-free rate puzzle. *J. Monetary Econ.* 24: 401–21
Xia Y. 2001. Learning about predictability: the effects of parameter uncertainty on dynamic asset allocation. *J. Financ.* 56:205–46.