On the Smooth Ambiguity Model: A Reply

Peter Klibanoff[*]   Massimo Marinacci[†]   Sujoy Mukerji[‡]

This version: September 6, 2009

Abstract

Epstein (2009) describes three Ellsberg-style thought experiments and argues that they pose difficulties for the smooth ambiguity model of decision making under uncertainty developed by Klibanoff, Marinacci and Mukerji (2005). We revisit these thought experiments and find, to the contrary, that they either point to strengths of the smooth ambiguity model compared to other models, such as the maxmin expected utility model (Gilboa and Schmeidler, 1989), or, in the case of one thought experiment, raise criticisms that apply equally to a broad range of current ambiguity models.

1   Introduction

Epstein (2009) describes three Ellsberg (1961)-style thought experiments and argues that they pose difficulties for the smooth ambiguity model of decision making under uncertainty developed by Klibanoff, Marinacci and Mukerji (2005, KMM). We revisit these thought experiments and find, to the contrary, that they either point to strengths of the smooth ambiguity model compared to other models, such as the maxmin expected utility (MEU) model (Gilboa and Schmeidler, 1989), or, in the case of one thought experiment, raise criticisms that apply equally to a broad range of current ambiguity models. We first discuss the second and third of these thought experiments, which are related, and close by discussing the first thought experiment.

To fix ideas and remind the reader, consider a separable metric state space Ω, a convex consequence space X, and the set Δ of all probability measures on Ω. The smooth ambiguity

---

[*]Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston, IL 60208, USA. E-mail: peterk@kellogg.northwestern.edu

[†]Collegio Carlo Alberto, Universita di Torino, Torino, Italy. E-mail: massimo.marinacci@carloalberto.org.

[‡]Department of Economics, University of Oxford, Oxford, U.K.. E-mail: sujoy.mukerji@economics.ox.ac.uk

model represents preferences over acts $f : \Omega \to X$ using the following functional:

$$V(f) = \int_\Delta \phi \left(\int_\Omega u\left(f\left(\omega\right)\right) dp\left(\omega\right) \right) d\mu(p), \tag{1.1}$$

where μ is a Borel probability measure on Δ, endowed with the vague topology.[1]

Similarly, the α-MEU model represents preferences over acts according to

$$U(f) = \alpha \inf_{p \in C} \int_\Omega u\left(f\left(\omega\right)\right) dp\left(\omega\right) + (1 - \alpha) \sup_{p \in C} \int_\Omega u\left(f\left(\omega\right)\right) dp\left(\omega\right), \tag{1.2}$$

where $\alpha \in [0, 1]$ is a weight and $C \subseteq \Delta$ is a set of probabilities. When $\alpha = 1$ we get back to the original MEU model.

2 Thought Experiment 2

Consider the second thought experiment proposed by Epstein. There is an urn containing 100 balls, of 4 different types: R_1, B_1, R_2, and B_2. A decision maker (DM) is told that $R_1 + B_1 = 50 = R_2 + B_2$, and that the relative proportions of (R_1, B_1) and (R_2, B_2) are determined independently. The DM considers bets with outcomes $c^* > c$ and the 50-50 lottery between them. For the sake of this discussion, let us assume that lotteries are evaluated according to an expected utility function u. This part of both the smooth ambiguity model and the MEU model is not central to the issues here. Normalize u so that $u(c^*) = 1$ and $u(c) = 0$. Then, $u(c^*, 1/2; c, 1/2) = u(c^*)/2 + u(c)/2 = 1/2$. With this notation, we can write the acts in question with utility payoffs as follows:

$$
\begin{array}{c|cccc}
& R_1 & B_1 & R_2 & B_2 \\
f_1 & 1 & 0 & 0 & 0 \\
f_2 & 0 & 0 & 1 & 0 \\
\frac{1}{2}f_1 + \frac{1}{2}f_2 & \frac{1}{2} & 0 & \frac{1}{2} & 0 \\
g_1 & \frac{1}{2} & \frac{1}{2} & 0 & 0 \\
g_2 & 0 & \frac{1}{2} & \frac{1}{2} & 0
\end{array}
\tag{2.1}
$$

Epstein claims, first, that given the symmetry in the problem, it is plausible that

$$f_1 \sim f_2.$$

---

[1] See KMM for details. Here, as in Epstein (2009), we use an Anscombe-Aumann version of the original KMM model.

We agree that symmetry makes this preference reasonable.

Second, he claims the preference

$$f_1 \sim f_2 \sim \frac{1}{2}f_1 + \frac{1}{2}f_2 \tag{2.2}$$

follows from the fact that the relative proportions of (R_1, B_1) and (R_2, B_2) are determined independently because, he claims, this should mean that f_1 and f_2 do not hedge one another. We will argue that *because* these relative proportions are independent, there is good reason why an ambiguity averse DM would see averaging f_1 and f_2 as desirable. In particular, we will show that it is only in very special circumstances that (2.2) would make sense.

Third, he claims that an ambiguity averse DM would have the preference

$$g_1 \succ g_2$$

because g_1 gives utility of $1/2$ on an unambiguous event with probability $1/2$, while g_2 pays off $1/2$ on a symmetric but ambiguous event. We agree with this third claim and the intuition for it.

Formally, his Proposition 3.1 shows that, in the context of the smooth ambiguity model, his three claims are incompatible. His argument proceeds roughly as follows: If, for a smooth ambiguity DM under the conditions in the proposition, one observes the preference $\frac{1}{2}f_1 + \frac{1}{2}f_2 \sim f_1 \sim f_2$, then there must be a range over which ϕ fails to be strictly concave (i.e., the DM is locally not ambiguity averse). On the other hand, $g_1 \succ g_2$ requires ϕ to be strictly concave (i.e., strictly ambiguity averse) in the relevant range. He interprets this result as showing that the smooth ambiguity model fails to capture intuitive choices in this thought experiment. We strongly disagree with this interpretation of the result, as we will argue that the second indifference in (2.2) is *not* intuitive for an ambiguity averse DM.

It useful to begin by writing down a simple, concrete set of probabilities that might plausibly be considered by the DM in the above example. In the context of the smooth ambiguity model, think of these as probabilities that are given uniform weight by the measure μ.

$$
\begin{array}{c|cccc}
 & R_1 & B_1 & R_2 & B_2 \\
\hline
p_1 & \frac{1}{10} & \frac{4}{10} & \frac{1}{10} & \frac{4}{10} \\
p_2 & \frac{1}{10} & \frac{4}{10} & \frac{4}{10} & \frac{1}{10} \\
p_3 & \frac{4}{10} & \frac{1}{10} & \frac{1}{10} & \frac{4}{10} \\
p_4 & \frac{4}{10} & \frac{1}{10} & \frac{4}{10} & \frac{1}{10} \\
\end{array}
\tag{2.3}
$$

Observe that these measures respect the given information, in that, for each probability p_i, $p_i(R_1 \cup B_1) = p_i(R_2 \cup B_2) = 1/2$. They respect the symmetry of the (lack of) information about the relative probabilities in each component in that for each marginal over (R_i, B_i) there is some measure with the corresponding marginal over (R_j, B_j), and these marginals are given identical weights by μ. Finally, these measures respect the believed independence of the two components in the sense that fixing a marginal over (R_i, B_i) leaves the same possibilities for the marginals over (R_j, B_j). In other words, informally, knowing the relative proportion of R_i to B_i places no restrictions on the relative proportion of R_j to B_j, and vice-versa. Note that this is the polar opposite of the situation where, whatever their value, the relative proportions of R_1 to B_1 and of R_2 to B_2 are known to be equal. This is the distinction between independently and indistinguishably distributed (IID) versus independently and identically distributed (iid) as discussed by, for example, Walley (1991) and Epstein and Schneider (2003).

Let us now examine the expected utilities generated by applying each of the measures p_i to the acts in the example. This will allow us to see the potential effect of the uncertainty over the probabilities on the DM. We begin with the acts g_1 and g_2. All four measures yield an expected utility of $1/4$ for g_1, reflecting the lack of ambiguity about the event $R_1 \cup B_1$. For g_2, however, the expected utilities vary, and are $5/20, 8/20, 2/20, 5/20$ respectively. Importantly, notice that the uniform average of these expected utilities is $1/4$ and any concave averaging will be less than $1/4$. Thus, any smooth ambiguity DM with a concave ϕ and a uniform μ will prefer g_1 to g_2. Intuitively, it is the dislike of the variation of expected utilities generated by the uncertainty about p that leads to this preference.

Consider next f_1, f_2, and $\frac{1}{2}f_1 + \frac{1}{2}f_2$. All of these are bets on ambiguous events, and the expected utilities generated by each of the p's are as follows:

	p_1	p_2	p_3	p_4
f_1	$\frac{1}{10}$	$\frac{1}{10}$	$\frac{2}{5}$	$\frac{2}{5}$
f_2	$\frac{1}{10}$	$\frac{2}{5}$	$\frac{1}{10}$	$\frac{2}{5}$
$\frac{1}{2}f_1 + \frac{1}{2}f_2$	$\frac{1}{10}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{2}{5}$

Notice that, again, respecting symmetry, the uniform averages of the expected utilities for each of these acts is $1/4$, thus an ambiguity neutral DM would quite reasonably be indifferent among all three. From symmetry alone, as reflected by the fact that the expected utilities for f_2 are a permutation of those for f_1, taking a common concave averaging would again yield the indifference between these first two acts. However, any concave averaging, because the distribution of expected utilities for both f_1 and f_2 is a mean-preserving spread

of that for $\frac{1}{2}f_1 + \frac{1}{2}f_2$, will lead to a higher average for $\frac{1}{2}f_1 + \frac{1}{2}f_2$ than for either f_1 or f_2. Thus, any smooth ambiguity DM with a strictly concave ϕ will have a strict preference for $\frac{1}{2}f_1 + \frac{1}{2}f_2$ over f_1 and f_2. This strict preference arises because the expected utilities vary across p_2 and p_3 under f_1 and f_2, but this variation is smoothed under $\frac{1}{2}f_1 + \frac{1}{2}f_2$. Observe that p_2 and p_3 correspond to the situation where there is one marginal over (R_1, B_1) and a different marginal over (R_2, B_2). In this sense, it is precisely *because* it is uncertain whether the two components are identical (thus leading p_2 and p_3 to be considered relevant) that the diversification provided by $\frac{1}{2}f_1 + \frac{1}{2}f_2$ is valuable. If, instead, the two components were known to be identical, even smooth ambiguity preferences would display the indifference in (2.2), as only p_1 and p_4 would be relevant. Notice, moreover, that, even in the indistinguishable but not necessarily identical case as in the example, an aggregation of the possible expected utilities taking into account only the minimum expected utility will lead to indifference between all three acts, as they all share the same minimum expected utility value, $1/10$, that is generated by p_1.

The observation above can be generalized beyond this specific set of priors and to α-MEU preferences. Let $\Omega = \{R_1, B_1, R_2, B_2\}$ and, given a set $C \subseteq \Delta$ of probabilities, set $\Gamma_i = \{p(R_i) : p \in C\}$. Consider the following properties:

(1) $\Gamma_1 = \Gamma_2$;

(2) $p(R_1 \cup B_1) = p(R_2 \cup B_2)$ for all $p \in C$;

(3) if $x \in \Gamma_1$ and $y \in \Gamma_2$, there is $p \in C$ such that $p(R_1) = x$ and $p(R_2) = y$.

For example, the set of priors in (2.3) satisfies (1)-(3), with $\Gamma_1 = \Gamma_2 = \{1/4, 4/10\}$. Property (1) reflects the symmetry across the components, (2) reflects the directly given information, and (3) seems a necessary condition for independence.

We can now state the following result, which is proved in the Appendix. The result references the condition

$$\mu(p \in C : p(R_1) \in B) = \mu(p \in C : p(R_2) \in B) \quad \text{for all Borel sets } B \subseteq [0, 1], \qquad (2.4)$$

which is meant to reflect the perceived symmetry across components.[2]

Proposition 2.1 *Suppose $C \subseteq \Delta$ is nonsingleton, closed, and satisfies properties (1)-(3). Then,*

---

[2] The sets $\{p \in C : p(R_i) \in B\}$ belong, for all Borel sets $B \subseteq [0, 1]$, to the Borel σ-algebra of Δ (see, e.g., Theorem 15.13 of Aliprantis and Border, 2006).

(i) *Any smooth ambiguity DM with ϕ strictly concave and μ with support C and such that condition (2.4) holds,[3] will have*

$$\frac{1}{2}f_1 + \frac{1}{2}f_2 \succ f_1 \sim f_2 \quad and \quad g_1 \succ g_2.$$

(ii) *Any α-MEU DM will have*

$$\frac{1}{2}f_1 + \frac{1}{2}f_2 \sim f_1 \sim f_2,$$

while $g_1 \succ g_2$ if and only if $\alpha > 1/2$.

Which of these preference patterns is more natural for an ambiguity averse DM? Return your attention to the original table describing the acts (2.1). The evaluation of act f_1 depends on the relative proportion of R_1 to B_1 but not on the relative proportion of R_2 to B_2. Similarly, the evaluation of f_2 depends on only the relative proportion of R_2 to B_2. In contrast, the evaluation of $\frac{1}{2}f_1 + \frac{1}{2}f_2$ depends on the relative proportions in both components, but has half the exposure to the uncertainty about the relative likelihood in each component compared to f_1 and f_2. Recall that the determination of the two relative proportions are viewed as independent. The act $\frac{1}{2}f_1 + \frac{1}{2}f_2$ thus diversifies the DM's exposure across the components: it provides a hedging of the two independent ambiguities in the same sense as diversifying across bets on independent risks provides a hedging of the risks. To a DM who is averse to ambiguity (i.e., to subjective uncertainty about relative likelihoods) because of the uncertainty it induces about expected utilities, such diversification is naturally valuable. Notice, however, that in this type of problem, preferences that ignore all except the minimum expected utility possibilities[4] will miss the diversification aspect of this situation, similar to an infinitely risk averse expected utility individual not valuing diversification across independent risks. This is extreme behavior. Models such as MEU and α-MEU force this extreme devaluation of diversification across these independent but not necessarily identical ambiguities. As we have argued, more moderate behavior, valuing the diversification, is natural for ambiguity averse DMs. The smooth ambiguity model delivers this moderate behavior, as it implies that such diversification is valued by almost all ambiguity averse DMs, though this value may vary in size as ambiguity attitude varies. Furthermore, when, *unlike* in this thought experiment, the two components are known to have identical marginals (and so are iid and not simply IID) and thus it is natural for the indifference in (2.2) to hold, it may be shown that the smooth ambiguity model indeed produces this.

To summarize our respective arguments regarding this interesting thought experiment

---

[3]Here the support of μ is defined as $\operatorname{supp}\mu = \bigcap \{D \text{ closed} : \mu(D) = 1\}$.

[4]Or a fixed weighted average of the minimum and maximum expected utility possibilities.

and its implications for the smooth ambiguity model: Epstein argues that $\frac{1}{2}f_1 + \frac{1}{2}f_2 \sim f_1 \sim f_2$ and $g_1 \succ g_2$ are natural for a strictly ambiguity averse DM, leading to a seeming inconsistency in the modelling of ambiguity attitude in the smooth ambiguity model through ϕ. We argue that for a strictly ambiguity averse DM, $\frac{1}{2}f_1 + \frac{1}{2}f_2 \succ f_1 \sim f_2$ is far more natural (or, at worst, no less natural) than $\frac{1}{2}f_1 + \frac{1}{2}f_2 \sim f_1 \sim f_2$. In this case, there is no conflict at all with $g_1 \succ g_2$, since both strict preferences are generated by a strictly concave ϕ in the smooth ambiguity model. Hence, we conclude, contrary to Epstein, that the choices in thought experiment 2 that are intuitive for an ambiguity averse DM are indeed captured within the smooth ambiguity model, whereas they are not captured (at least with symmetry and independence modelled as above) by the MEU (or α-MEU) model.[5] Beyond the specific issue of compatibility with the smooth ambiguity model, this discussion and thought experiment highlights a point we feel is fundamental in thinking about ambiguity aversion – hedging across independent sources of ambiguity makes a lot of sense.

3   Thought Experiment 3

Epstein poses this example to claim that "calibration of ambiguity aversion within the KMM model, or the transferability of the function ϕ across settings, is, in general, inconsistent with intuitive choices." He first presents Urn I, which is a situation identical to thought experiment 2, except he introduces some information about the composition of the second component designed to break the symmetry between components. Specifically, at least 10 of the 50 balls are known to be Red in the second component. Then he restores indifference between a bet, f_1^I, on Red in the first component and a bet, f_2^I, on Red in the second component by raising the payoff on Red for the bet on the first component until the indifference, $f_1^I \sim f_2^I$, is attained. Next, he asserts, just as in thought experiment 2, that the independence of the two components leads to $\frac{1}{2}f_1^I + \frac{1}{2}f_2^I \sim f_1^I \sim f_2^I$. He also presents an Urn II, differing from Urn I only in that, instead of being told that at least 10 of the 50 balls are Red in the second component, we are told that exactly half of the balls in the urn are Red (i.e., $R_1 + R_2 = 50$). With this new information, a 50-50 mixture of a bet on Red in the first component, f_1^{II}, and the same stakes bet on Red in the second component, f_2^{II}, will create an unambiguous act, $\frac{1}{2}f_1^{II} + \frac{1}{2}f_2^{II}$. Thus, he argues and we agree, ambiguity aversion suggests

[5]By no means do we claim that the smooth ambiguity model (and its close relatives Nau, 2006, Ergin and Gul, 2009, Neilson, 2009, and Seo, forthcoming) is the only model capturing these intuitive choices. Many other models in the ambiguity aversion literature – e.g., invariant biseparable preferences (Ghirardato, Maccheroni and Marinacci, 2004, and Amarante, 2009), variational preferences (Maccheroni, Marinacci and Rustichini, 2006), and vector expected utility preferences (Siniscalchi, 2009) – have cases compatible with the choices that we claim are intuitive.

$\frac{1}{2}f_1^{II} + \frac{1}{2}f_2^{II} \succ f_1^{II} \sim f_2^{II}$. In Proposition 4.1, he shows that given $\frac{1}{2}f_1^I + \frac{1}{2}f_2^I \sim f_1^I \sim f_2^I$ (and given $g_1 \succ g_2$, so that the situation where f_1^I and f_2^I yield the same expected utility for all probabilities considered by the DM is ruled out), ϕ must be affine over a given range, and this implies that, for some level of payoffs, the intuitive strict preference $\frac{1}{2}f_1^{II} + \frac{1}{2}f_2^{II} \succ f_1^{II} \sim f_2^{II}$ in Urn II is not possible.

As this exposition hopefully makes clear, this example is simply an elaboration of thought experiment 2, and the driving force behind the conclusion is exactly the same as there: assuming $\frac{1}{2}f_1^I + \frac{1}{2}f_2^I \sim f_1^I \sim f_2^I$ forces ϕ to be affine over some range (thus generating ambiguity neutral behavior over that range) and leads the DM to assign no value to the elimination of ambiguity provided by the act $\frac{1}{2}f_1^{II} + \frac{1}{2}f_2^{II}$ over that range in Urn II. Therefore, exactly the same argument as we presented above, calling into question Epstein's assertion that there is nothing gained by spreading bets across independent sources of ambiguity, applies here. As we have argued, this is an unwarranted assertion: in fact, $\frac{1}{2}f_1^I + \frac{1}{2}f_2^I \succ f_1^I \sim f_2^I$ is quite natural for an ambiguity averse DM. The smooth ambiguity model with a single, strictly concave ϕ is perfectly compatible with this latter assumption and the remaining preferences in Urns I and II. Hence, we conclude, contrary to Epstein, that thought experiment 3 does *not* demonstrate any problem with the transferability of the function ϕ across the two urns.

While we think thought experiment 3 does not shed light on transferability issues, we do agree that such issues are subtle and require careful thought. In particular, there are important limitations related to Savage's (1972) "small worlds" discussion which are relevant when applying a wide range of models (including not only the smooth ambiguity model, but also subjective expected utility, MEU, and many others). In that discussion, Savage emphasizes the uncertain or incompletely observed nature of consequences in any real application of utility theory and summarizes as follows (p. 84): "Indeed, in the final analysis, a consequence is an idealization that can perhaps never be well approximated." This limits greatly the extent to which taste parameters, such as risk and ambiguity attitudes, can be measured in one context and that measurement directly transferred to a very different context. Since two contexts may "complete" the consequences in radically different ways, acts in the two settings which give the same "nominal" consequences might appear to be the same, but in reality differ. Therefore, even when underlying tastes are stable across the settings, it would be inappropriate to apply the same functional representation of those tastes to these "nominal" acts. A simple example in the context of expected utility theory is the following: Suppose we observe an individual's choices among monetary gambles in a laboratory experiment. Further, suppose that the individual chooses according to expected utility with a utility function that treats monetary outcome \m as the "consequence" $m + c$, $u(m + c)$, where c here is meant to stand for other things the individual cares about that do not vary

with the outcome of the monetary gambles and are unobserved by the experimenter. From this choice data, the experimenter infers a utility function over money, $\hat{u}(m)$. Despite the fact that c is unobservable (and, in fact, is ignored by the experimenter), $\hat{u}$ will accurately represent the risk preferences of the individual over monetary gambles *in this context* (i.e., the risk preferences represented by $u(\cdot + c)$). To emphasize this last qualification, suppose we observe the same individual again choosing among monetary gambles, but now in a different (for example, more "real-life" and/or later in time) setting. Suppose that the individual now treats monetary outcome \$$m$ as the "consequence" $m + d$ where $d \neq c$. This is meant as a simple device to reflect the change in the individual's circumstances or consumption opportunities or related matters. Assume choices are still determined according to the expectation of $u(m + d)$. In general, the utility function over money, $\tilde{u}(m)$, inferred from the second set of choices will not be equivalent to the $\hat{u}(m)$ elicited earlier even though the utility u defined over the complete consequences *has remained unchanged*. Thus, even though risk attitudes over consequences are stable across the two settings (by assumption), using the risk attitude toward monetary gambles measured in the first setting to predict choices over monetary gambles in the second setting may fail.

Despite this, within settings, separation of tastes, such as risk or ambiguity attitudes, from beliefs is important and useful. For instance, the risk aversion calibrated within a setting can be used to predict further choices in that same setting. In fact, this "internal calibration" is the nature of many common economic applications e.g., the equity premium/risk-free rate puzzle. The equity premium puzzle may be posed purely in terms of internal calibration as follows: Take a "fruit-tree economy" à la Lucas (1978). In such an economy, given an expected utility maximizing representative agent whose belief is calibrated to salient aspects of the historical/empirical distribution of (consumption) growth (e.g., matching various moments), it is impossible to calibrate the agent's risk aversion parameter such that both the equilibrium return on the riskfree asset and that on the riskless asset match those in the observed data. If the risk aversion parameter is set such that the equilibrium risk free rate matches data, the return on the risky asset is far too low; if the risk aversion parameter is set such that the return on the risky asset matches the data, the implied equilibrium return on the riskfree asset is far too high. Note, to articulate the puzzle, we needed a decision model with a conceptual and applicable separation of risk and risk attitude, so that within the general equilibrium model risk (as modelled by the belief) and risk attitude could be calibrated independently. Just as important to note, the puzzle is framed as a problem of inability to calibrate the risk aversion parameter in an internally consistent manner. In particular, the key issue is calibration within a single domain, and not calibration of risk aversion across domains.

Similarly, a key issue in a number of recent papers on ambiguity and asset pricing – e.g., Chen, Ju, and Miao (2009), Hansen (2007), and Hansen and Sargent (2008) – is whether replacing the expected utility agent by (variations of) a smooth ambiguity agent and calibrating (ambiguous) beliefs and ambiguity attitude (holding risk aversion fixed at an externally calibrated level) within the single domain allows one to better match the data on the riskfree and risky returns simultaneously. As many of their findings show, it is possible to obtain an internally consistent calibration such that both rates of return are (largely) consistent with data. Cross-domain transferability is no more problematic in this application of the smooth ambiguity model than it is for the standard expected utility model. In any case, it is not Epstein's thought experiment, but rather Savage's "small worlds" discussion that seems to us most relevant to these issues.

4   Thought Experiment 1

The experiment takes Ellsberg's (1961) 3-color problem and adds a "second-order" urn, also containing 90 total balls. A draw from this second-order urn will determine the composition of the Ellsberg 3-color urn. Specifically, 30 of the 90 balls in the second-order urn result in the odds $p_1 = (1/3, 1/3, 1/3)$ in the Ellsberg urn, while the remaining balls in the second-order urn are divided in an unspecified manner between those that give odds $p_2 = (1/3, 0, 2/3)$ and $p_3 = (1/3, 2/3, 0)$ in the Ellsberg urn. In addition to the usual bets on the color of the ball drawn from the Ellsberg urn, Epstein considers bets on the composition of the Ellsberg urn (equivalently, bets on the type of ball drawn from the second-order urn). He argues that the standard Ellsberg choices over bets about the color drawn from the Ellsberg urn would imply Ellsberg choices over bets on the composition of the Ellsberg urn. This is his "intuitive hypothesis." Furthermore, he claims that the MEU model satisfies the intuitive hypothesis while the smooth ambiguity model does not. Below, we refute the first part of this claim and offer remarks about the intuitive hypothesis more generally.

First, we show that the MEU model does not deliver choices respecting the "intuitive hypothesis." Taking the relevant state space, as Epstein does, to be $\{R, G, B\}$, his hypothesis in this thought experiment is that the preference over acts based on the Ellsberg urn – i.e., $f_1 \succ f_2$ and $f_4 \succ f_3$ as stated in (2.1) of Epstein (2009) – implies Ellsbergian preferences over second-order acts ("bets on the composition"), i.e., $F_1 \succ F_2$ and $F_4 \succ F_3$ as stated in (2.2) of Epstein (2009).[6] To extend the MEU model to preferences over second-order acts,

---

[6]Note that, literally, the second-order acts F_1, F_2, F_3, and F_4 are not completely specified, as the table does not specify what payoff is received if an urn composition other than p_1, p_2, p_3 is realized. From the verbal description the natural way to complete the specification is to assign 0 if any other composition occurs, and so we assume this. Nothing material would change in our results if this payoff were different than 0 as

Epstein proposes the following representing functional (see his (2.4)),

$$U^2(F) = \inf_{p \in C} u(F(p)) \tag{4.1}$$

wherein a second-order act, F, is evaluated using the minimum utility that it gives under any of the probability measures in the set C used in the usual MEU representation of preferences over acts (see (1.2) with $\alpha = 1$ or his (2.3)). Then, in the middle of p. 7 he states "It is obvious that the multiple-priors model satisfies the intuitive hypothesis in our experiment...." However, using (1.2) and (4.1) to evaluate acts and second order acts, respectively, leads to violation of his hypothesis.

Observation 4.1 *Fix any set C such that $\{p_1, p_2, p_3\} \subseteq C \subseteq \Delta\{R, G, B\}$, where $p(R) = 1/3$ for all $p \in C$. According to (1.2) with $\alpha = 1$, $f_1 \succ f_2$ and $f_4 \succ f_3$, while according to (4.1), $F_1 \sim F_2 \sim F_3 \sim F_4$.*

Proof. By (1.2) with $\alpha = 1$, $U(f_1) = U(f_3) = u(100)/3$, $U(f_2) = u(0)$, and $U(f_4) = 2u(100)/3$, and so $f_1 \succ f_2$ and $f_4 \succ f_3$. By (4.1), $U(F_1) = U(F_2) = U(F_3) = U(F_4) = u(0)$ since each of these second-order acts has a common minimum payoff of 0. ∎

The only way to reach a different conclusion, it seems to us, would be to change the state space to $\{R, G, B\} \times \{p_1, p_2, p_3\}$. But then, it is easy to satisfy Epstein's hypothesis with either the smooth ambiguity model or the multiple priors model, as in each, the additional probabilistic information about the distribution over $\{p_1, p_2, p_3\}$ (specifically, that p_1 occurs with probability 1/3) may be incorporated directly into the set C in the same way that it was concerning $\{R, G, B\}$. On p. 6, Epstein argues that moving to the expanded state space does not solve the problem, as examples involving bets on probabilities over $\{R, G, B\} \times \{p_1, p_2, p_3\}$ could then be constructed. Our point here is that the same "problem," illustrated in Observation 4.1, would exist for the MEU model one level up as well. We believe that whatever message is conveyed by this thought experiment, if one is to interpret it as a critique of the smooth ambiguity model then, by our arguments above, one should equally interpret it as a critique of the MEU model.

Our second remark concerns the more general point of this thought experiment. Namely, strong asymmetries in the quality of information about the likelihood of different possible probabilities make violations of the sure-thing principle more plausible when choosing among bets on probability distributions. In the thought experiment, this asymmetry is operational-ized by specifying the likelihood of probability p_1 precisely (30 of 90 balls in the second-order urn) while leaving the relative likelihoods of p_2 and p_3 unspecified. We largely agree with

_____
long as it was the same for all four second-order acts.

this point, though, based on our discussion above, we disagree that this is a critique unique to the smooth ambiguity model. We think this indicates a more general limitation present in many models of decision making under ambiguity: for events over which (first-order) acts are defined, variation in the extent of ambiguity is allowed and captured, while events at higher levels are modeled as having a uniform degree of ambiguity. To understand the nature of this limitation, let us focus on the smooth ambiguity model.

Intuitively, ambiguity averse DMs prefer acts whose evaluation is more robust to possible variation in probabilities. In KMM that is translated as an aversion to mean preserving spreads in the induced distribution of expected utilities. This is shown to be equivalent to concavity of ϕ and to the DM being more averse to the subjective uncertainty about priors than he is to the risk in lotteries. It is as if we imagine an ambiguity averse DM to be thinking as follows. "My best guess of the chance that the return distribution is 'p' is 20%. However, this is based on 'softer' information than knowing that the chance of a particular outcome in an objective lottery is 20%. Hence, I would like to behave with more caution with respect to the former risk."

This caution translates into behavior in the smooth ambiguity model in the following way. When the DM is strictly ambiguity averse, a bet that the return distribution is 'p' is valued less than a same-stakes lottery with a 20% chance of winning, while, at the same time a bet *against* the return distribution being 'p' is also valued less than a same-stakes lottery with an 80% chance of winning. In general, strict ambiguity aversion will thus make it impossible to calibrate the likelihood of second-order events to lottery probabilities. In this sense, analogous to the definition of ambiguous events given in Section 4 of KMM, these second-order events are ambiguous. In fact, strict ambiguity aversion combined with non-degenerate beliefs over the return distribution implies that *all* (non-empty and non-universal) second-order events are ambiguous. Similarly, when beliefs over the return distribution are degenerate, all second-order events are unambiguous. It is in this sense that a uniformity is imposed on the ambiguity of events at the second-order level, that is most certainly not imposed on events at the level of the space on which (first-order) acts are defined. This is why a situation as in the thought experiment, where the second-order event 'p_1' is plausibly unambiguous while the second-order events 'p_2' and 'p_3' are plausibly ambiguous, is not able to be handled by the smooth ambiguity model.

A related limitation of the model, not directly evident in the thought experiment, is that, while second-order events are treated as more ambiguous than lottery events, to the extent that third-order, fourth-order or higher level events are considered, the model treats them as second-order events. This is a second type of uniformity imposed by the smooth ambiguity

12

model.[7]

Both of these uniformities (shared also by the MEU model and others) are implications that undoubtedly conflict with some plausible behavior, but are the result of exactly the kind of pragmatic modeling choices that are often useful in building tractable and insightful models. They allow us to focus attention of the domain of primary interest, preferences over (first-order) acts, and the (allowed) non-uniformity of information quality relevant to decisions on that domain.

In sum, this thought experiment raises issues of variation in the quality of information about higher-order events and the DM's reaction to this variation. The smooth ambiguity model (as well as MEU and many other models) does not naturally incorporate such variation. We have argued that such a limitation should be viewed as a pragmatic modeling choice and is of little importance in many applications. Furthermore, to the extent that, in a particular application, this variation up to a given order of events is deemed important, there is, as Epstein and we have noted, an easy way, both for MEU and for the smooth ambiguity model, to capture it. Simply expand the state space on which acts are defined to incorporate events up to that order. In the thought experiment, this corresponds to incorporating the true probability law, by using $\{R, G, B\} \times \{p_1, p_2, p_3\}$ as the state space.

5 Appendix: Proof of Proposition 2.1

Observe that, under (1) and (2), C nonsingleton implies Γ_i nonsingleton for $i = 1, 2$. By (3), this implies that there is $p \in C$ such that $p(R_1) \neq p(R_2)$.

(i) Suppose $\operatorname{supp} \mu = C$ and $\mu(p \in C : p(R_1) \in B) = \mu(p \in C : p(R_2) \in B)$ for all Borel sets B in $\mathbb{R}$. Since ϕ is strictly increasing, by (1) we have $\{(\phi \circ p)(R_1) : p \in C\} = \{(\phi \circ p)(R_2) : p \in C\}$, and so $\int_\Delta (\phi \circ p)(R_1) \, d\mu(p) = \int_\Delta (\phi \circ p)(R_2) \, d\mu(p)$ because of the assumption on μ. Hence, $f_1 \sim f_2$. On the other hand,

$$\phi\left(\frac{1}{2}p(R_1) + \frac{1}{2}p(R_2)\right) \geq \frac{1}{2}(\phi \circ p)(R_1) + \frac{1}{2}(\phi \circ p)(R_2), \qquad \forall p \in \operatorname{supp} \mu,$$

with strict inequality if $p(R_1) \neq p(R_2)$.

Claim There is a Borel set $A \subseteq \operatorname{supp} \mu$, with $\mu(A) > 0$, such that $p(R_1) \neq p(R_2)$ for all $p \in A$.

---

[7]In the terminology of the source-preference literature (see, e.g., Tversky and Fox, 1995, Chew and Sagi, 2008, and Wakker, 2008) the space of second-order and higher events is a single source of uncertainty as is the space of lotteries. For comparisons within a single source, preferences are expected utility, but not for comparisons across sources, such as comparisons of bets on second-order events with lotteries.

Proof of the Claim There is $\bar{p} \in \operatorname{supp}\mu$ such that $\bar{p}(R_1) \neq \bar{p}(R_2)$. Suppose first that $\bar{p}$ is an isolated point in $\operatorname{supp}\mu$. Then, $\mu(\bar{p}) > 0$ and the claim trivially holds. Suppose that $\bar{p}$ is not an isolated point in $\operatorname{supp}\mu$. Then, $B_\varepsilon(\bar{p}) \cap \operatorname{supp}\mu \neq \emptyset$ for every neighborhood $B_\varepsilon(\bar{p})$ of $\bar{p}$. Since $\bar{p}(R_1) \neq \bar{p}(R_2)$, by taking ε small enough there is $B_\varepsilon(\bar{p})$ such that $p(R_1) \neq p(R_2)$ for all $p \in B_\varepsilon(\bar{p})$. By setting $A = B_\varepsilon(\bar{p}) \cap \operatorname{supp}\mu$, this proves the claim since $\mu(A) > 0$ because $B_\varepsilon(\bar{p}) \cap \operatorname{supp}\mu \neq \emptyset$. For, if $\mu(A) = 0$, then $\mu(B_\varepsilon(\bar{p})) = \mu(A) + \mu(B_\varepsilon(\bar{p}) \cap (\operatorname{supp}\mu)^c) = 0$, and so $\operatorname{supp}\mu \subseteq B_\varepsilon(\bar{p})^c$, a contradiction (see Aliprantis and Border, 2006, p. 442). △

The Claim implies

$$\int \phi\left(\frac{1}{2}p(R_1) + \frac{1}{2}p(R_2)\right) d\mu(p) > \frac{1}{2}\int (\phi \circ p)(R_1) d\mu(p) + \frac{1}{2}\int (\phi \circ p)(R_2) d\mu(p),$$

that is, $\frac{1}{2}f_1 + \frac{1}{2}f_2 \succ f_1 \sim f_2$.

We have $\phi(p(R_1 \cup B_1)/2) = \phi(1/4)$ for all $p \in \operatorname{supp}\mu$. Moreover, $\phi(p(B_1 \cup R_2)/2) = \phi(1/4 + (p(R_2) - p(R_1))/2)$. Define $\gamma : \Delta \to \mathbb{R}$ by $\gamma(p) = 1/4 + (p(R_2) - p(R_1))/2$. Since $\gamma(p) \neq 1/4$ for all $p \in A$, by the Jensen inequality we have

$$\int (\phi \circ \gamma)(p) d\mu(p) < \phi\left(\int \gamma(p) d\mu(p)\right) = \phi\left(\int \left(\frac{1}{4} + \frac{1}{2}(p(R_2) - p(R_1))\right) d\mu(p)\right) = \phi\left(\frac{1}{4}\right),$$

that is, $g_1 \succ g_2$.

(ii) By (1)-(3), $\max_{p \in C} p(R_1) = \max_{p \in C} p(R_2)$ and $\min_{p \in C} p(R_1) = \min_{p \in C} p(R_2)$, as well as

$$\begin{aligned}
\max_{p \in C}\left(\frac{1}{2}p(R_1) + \frac{1}{2}p(R_2)\right) &= \frac{1}{2}\max_{p \in C} p(R_1) + \frac{1}{2}\max_{p \in C} p(R_2) \\
\min_{p \in C}\left(\frac{1}{2}p(R_1) + \frac{1}{2}p(R_2)\right) &= \frac{1}{2}\min_{p \in C} p(R_1) + \frac{1}{2}\min_{p \in C} p(R_2)
\end{aligned}$$

Hence, $f_1 \sim f_2$. From $\min_{p \in C}(p(R_2) - p(R_1)) = -\max_{p \in C}(p(R_2) - p(R_1)) < 0$, it follows

$$\begin{aligned}
&\alpha \min_{p \in C}\left(\frac{1}{2}p(B_1) + \frac{1}{2}p(R_2)\right) + (1-\alpha)\max_{p \in C}\left(\frac{1}{2}p(B_1) + \frac{1}{2}p(R_2)\right) \\
&= \frac{1}{4} + \frac{1-2\alpha}{2}\max_{p \in C}(p(R_2) - p(R_1)),
\end{aligned}$$

and so $g_1 \succ g_2$ if and only if $1/4 > 1/4 + (1/2 - \alpha)\max_{p \in C}(p(R_2) - p(R_1))$, i.e., if and only if $\alpha > 1/2$. ∎

14

References

[1] C. D. ALIPRANTIS AND K. C. BORDER, *Infinite dimensional analysis* (2006), 3rd ed., Springer Verlag, Berlin.

[2] M. AMARANTE, Foundations of neo-Bayesian statistics, *Journal of Economic Theory* **144** (2009), 2146-2173.

[3] H. CHEN, N. JU, J. MIAO, Dynamic asset allocation with ambiguous return predictability, mimeo, *Boston University*, 2009.

[4] S.H. CHEW, J. SAGI, Small worlds: Modeling attitudes toward sources of uncertainty, *Journal of Economic Theory* **139** (2008), 1-24.

[5] D. ELLSBERG, Risk, ambiguity and the Savage axioms, *Quarterly Journal of Economics* **75** (1961), 643-669.

[6] L.G. EPSTEIN, Three paradoxes for the 'smooth ambiguity' model of preference, mimeo, *Boston University* (2009).

[7] L.G. EPSTEIN, M. SCHNEIDER, IID: Independently and indistinguishably distributed, *Journal of Economic Theory* **113** (2003), 32-50.

[8] H. ERGIN, F. GUL, A subjective theory of compound lotteries, *Journal of Economic Theory* **144** (2009), 899-929.

[9] P. GHIRARDATO, F. MACCHERONI, M. MARINACCI, Differentiating ambiguity and ambiguity attitude, *Journal of Economic Theory* **118** (2004), 133-173.

[10] I. GILBOA, D. SCHMEIDLER, Maxmin expected utility with non-unique prior, *Journal of Mathematical Economics* **18** (1989), 141-153.

[11] L.P. HANSEN, Beliefs, doubts and learning: The valuation of macroeconomic risk, *American Economic Review* **97** (2007), 1-30. 1144-1152.

[12] L.P. HANSEN, T.J. SARGENT, Fragile beliefs and the price of model uncertainty, mimeo, *University of Chicago*, 2008.

[13] N. JU, J. MIAO, Ambiguity, learning, and asset returns, mimeo, *Boston University*, 2009.

[14] P. KLIBANOFF, M. MARINACCI, S. MUKERJI, A smooth model of decision making under ambiguity, *Econometrica* **73** (2005), 1849-1892.

[15] R.E. Lucas, Asset Prices in an Exchange Economy, *Econometrica* **46** (1978), 1429-1445.

[16] F. Maccheroni, M. Marinacci, A. Rustichini, Ambiguity aversion, robustness, and the variational representation of preferences, *Econometrica* **74** (2006), 1447-1498.

[17] R. Nau, Uncertainty aversion with second-order utilities and probabilities, *Management Science* **52** (2006), 136-145.

[18] M. Siniscalchi, Vector expected utility and attitudes toward variation, *Econometrica* **77** (2009), 801-855.

[19] W.S. Neilson, A simplified axiomatic approach to ambiguity aversion, mimeo, *University of Tennessee* (2009).

[20] L.J. Savage, *The foundations of statistics* (1972), 2nd ed., Dover, New York.

[21] D. Schmeidler, Subjective probability and expected utility without additivity, *Econometrica* **57** (1989), 571-587.

[22] K. Seo, Ambiguity and second-order belief, *Econometrica, forthcoming.*

[23] A. Tversky, C. Fox, Weighing risk and uncertainty, *Psychological Review* **102** (1995), 269-283.

[24] P.P. Wakker, Uncertainty, in L. Blume and S. N. Durlauf (eds.), *The new Palgrave: A dictionary of economics*, (2008), 6780-6791, The MacMillan Press, London.

[25] P. Walley, *Statistical reasoning with imprecise probabilities* (1991), Chapman and Hall, London.