PIER Working Paper 15-009

"Optimism and Pessimism with Expected Utility
Fifth Version"

by

David Dillenberger, Andrew Postlewaite, and Kareen Rozen

http://ssrn.com/abstract=2568696

Optimism and Pessimism with Expected Utility[*]

David Dillenberger
University of Pennsylvania

Andrew Postlewaite[†]
University of Pennsylvania

Kareen Rozen[‡]
Yale University

February 2015

Abstract

Maximizing subjective expected utility is the classic model of decision making under uncertainty. Savage (1954) provides axioms on preference over acts that are equivalent to the existence of a subjective expected utility representation, and further establishes that such a representation is essentially unique. We show that there is a continuum of other "expected utility" representations in which the probability distributions over states used to evaluate acts depend on the set of possible outcomes of the act and suggest that these alternate representations can capture pessimism or optimism. We then extend the DM's preferences to be defined over both subjective acts and objective lotteries, allowing for source-dependent preferences. Our result permits modeling ambiguity aversion in Ellsberg's two-urn experiment using a single utility function and pessimistic probability assessments over prizes for lotteries and acts, while maintaining the axioms of Savage and von Neumann-Morganstern on the appropriate domains.

If one has really technically penetrated a subject, things that previously seemed in complete contrast, might be purely mathematical transformations of each other. —John von Neumann (1955, p. 496).

1 Introduction

Consider a decision maker (DM) who is faced with gambles on whether it will rain more in Northern Ghana (N) than in Southern Ghana (S) tomorrow. He is told that if the outcome is (N) he will get \$100 and if (S) he will also get \$100. When asked what he thinks the probability of N is, he responds .5. He is then told about another gamble in which the outcome for S is unchanged but the outcome for N is increased to \$1000, and is asked what he thinks the probability of N is now. He responds that he thinks the probability of N is now .4. When asked how he can think the probability of N can differ across the two gambles when it is the same event, DM simply says that random outcomes tend to come out badly for him. After being offered a third gamble that gives \$100 for S and \$10,000 for N, he says that faced with that gamble, he thinks the probability of N is .2.

When faced with a choice between any two gambles, each of which specifies the amount received conditional on the realized state, DM says that he maximizes expected utility. He has a utility function over money, and for any two gambles (x_1, x_2) and (y_1, y_2), he will have two probability distributions over the states, $p(x_1, x_2)$ and $p(y_1, y_2)$. DM's probability assessments reflect his belief that luck is not on his side. For each gamble he computes its expected utility under the associated probability, and then chooses the gamble with the higher expected utility.

Confronted with such a DM, one might well judge him irrational. But would that judgment change if one discovered that the DM's revealed preferences satisfy Savage's axioms? We show below that for any preferences over acts that satisfy Savage's axioms, there will be representations of those preferences as described in the paragraph above: there will be a utility function over outcomes and, for any act, a probability distribution over states that depends on the payoffs the act generates, with preferences given by expected utility. Furthermore, the probability distribution depends on the payoffs as in the example above: the probability of the state with the good outcome is smaller than the Savage probability, and it decreases when the good outcome is replaced by an even better outcome. We suggest that a DM who describes his decision-making process as above can be thought of as pessimistic. Similarly, in addition to the multitude of pessimistic representations of preferences that satisfy Savage's axioms, there is a continuum of "optimistic" representations.

We may still want to characterize the DM above as being irrational, but notice that we cannot make that determination on the basis of his choices: his preferences over acts are the same as those of a person who uses an analogous decision process using the Savage representation utility function and associated "standard" probability distribution. Any distinction between the rationality of

the Savage representation and the alternative representation must be made on the basis of the underlying process by which the DM makes decisions and not only on the decisions themselves.

We then extend the DM's preferences to be defined over both subjective acts and objective lotteries. We show that in this extended domain, we can address Ellsberg (1961)'s two-urn experiment using standard expected utility in the objective world and "pessimistic," stake-dependent expected utility (in the sense above) in the subjective world, while applying the *same* utility over prizes in both domains. This relates to the literature on source-dependent preferences (Chew and Sagi (2008), among others), which also addresses Ellsberg's experiment without relaxing the appealing axioms of Savage and vNM on the respective domains, but has been criticized for capturing ambiguity attitudes by a source-dependent utility function over prizes rather than different probability assessments (see, for example, Wakker (2010, p. 337)). Indeed, we believe the latter would be more natural in this context since, intuitively, it is the probability assessments, not the prizes, that change when moving to the uncertainty domain.

A consequence of the multiplicity of alternative representations of preferences that satisfy Savage's axioms is that existing analyses of agents' market behavior in the face of uncertainty have a broader interpretation than would appear at first glance. It might have seemed as though a model that assumed that agents maximized expected utility with "standard" stake-independent probabilities had little or nothing to say about how a pessimistic agent, in the sense we outlined above, would behave. But since the choices of the pessimistic agent depend only on the underlying preferences, if those underlying preferences have in addition a representation with stake-independent probabilities, then the pessimistic agent's behavior conforms precisely to any predictions that come from the standard model despite his unorthodox mental processing. In short, agents who are pessimistic or optimistic in our sense are observationally indistinguishable from the "standard" expected utility maximizing agent. This connection is, of course, based on our cognitive notion of pessimism, which differs from other notions of pessimism that are based solely on observed choice behavior, as discussed in Section 4.2. But it does suggest that one need not necessarily modify the standard model to include psychologically plausible decision processes that differ from that associated with the Savage representation.[1]

It is useful to distinguish between a *utility representation* (or *model*), which is a construct for imagining how a DM makes decisions, and *choice behavior*, which is the observable data. The standard point of view is that the representation is nothing more than an analytically convenient device to model a DM's choices. In this approach, termed *paramorphic* by Wakker (2010), the representation does *not* suggest that a DM uses the utility function and a probability distribution

---

[1]Hey (1984), for example, introduces a notion of pessimism and optimism very similar to our own: an optimist (pessimist) revises up (down) the probabilities of favorable events and revises down (up) the probabilities of unfavorable events. Hey incorporates consequence-dependent probabilities in a Savage-like representation, which can generate behavioral patterns that are inconsistent with expected utility because additional restrictions are not placed on the distorted probabilities. The notion that optimism and pessimism are inconsistent with Savage's axioms is implicit in his analysis, whereas our paper suggests that this is not necessarily the case.

to make choices. An alternative approach is that the models we employ should not only capture the choices agents make, but should match the underlying processes in making decisions. Wakker (2010) lays out an argument for this approach, which he terms *homeomorphic*. In his words, "we want the theoretical parameters in the model to have plausible psychological interpretations." This stance is also common in the behavioral economics literature, where mental processes and psychological plausibility are of particular interest.[2,3]

It is sometimes productive to take the elements of the representation as actual entities in themselves. Consider a situation in which a DM may have little or no information about the relative likelihoods of outcomes associated with different choices she confronts. An (unbiased) expert who is informed about those likelihoods could determine which of the choices is best if he knew the DM's utility function. Through a sequence of questions about choices in a framework that the DM understands, the expert can, in principle, elicit the utility function, which can then be combined with the expert's knowledge about the probabilities associated with the choices in the problem at hand in order to make recommendations. Wakker (2008, 2010) and Karni (2009) treat problems of this type in the context of medical decision making. Under this point of view, it may be important to understand which representation is being elicited. If a DM had stake-dependent pessimistic beliefs but was assumed to have a "standard" Savage representation, the elicited utility function would exhibit greater risk aversion than the true utility function. Analogously, for an optimistic DM, the elicited utility function would exhibit less risk aversion than her true utility function.

The remainder of this paper is organized as follows. We lay out the model in Section 2 and demonstrate how pessimistic and optimistic representations can be constructed. In Section 3 we study the extension of the DM's preferences to both subjective acts and objective lotteries. Section 4 discusses related work.

---

[2]A similar discussion appears in Karni (2011). Karni distinguishes between the definitional meaning of subjective probabilities, according to which subjective probabilities define the DM's degree of belief regarding the likelihood of events, and the measurement meaning, according to which subjective probabilities measure, rather than define, the DM's beliefs. That is, the DM's beliefs are cognitive phenomena that directly affect the decision-making process.

[3]As Dekel and Lipman (2010) note, a utility representation is, at minimum, useful for organizing our thoughts around the elements of that representation (e.g., in terms of probabilities, utilities, and expectations). They further argue that the "story" of a model is relevant and may provide a reason for preferring one model to the other, even if the two models predict the same choices. Saying that, Dekel and Lipman emphasize that while the story's plausibility (or lack thereof) may affect our confidence in the predictions of the model, it cannot refute or confirm those predictions; and that even if the story suggested by the representation is known to be false, it may still be valuable to our reasoning process.

2 Optimism, pessimism, and stake-dependent probabilities

2.1 Two states of nature

There are two states of nature, s_1 and s_2. Let $X \subset \mathbb{R}$ be an interval of monetary prizes. Consider a DM whose preferences over the set of (Savage) acts satisfy Savage's axioms, and who prefers more money to less.[4] Formally, an act is a function $l : \{s_1, s_2\} \to X$. For notational convenience, in the text we simply denote an act by an ordered pair of state contingent payoffs, $x = (x_1, x_2)$, where x_i is the payoff received in state i. Let $v(x) = p_1 u(x_1) + p_2 u(x_2)$ represent the DM's preferences over acts. Here $p = (p_1, p_2)$ is the subjective, stake-independent probability distribution.

We now consider a different representation of the same preferences, in which the probability distribution is *stake-dependent*: that is, the probability assigned to each state i is $P_i(x; p)$. We look for a representation $\hat{v}$ of the form

$$\hat{v}(x) = P_1(x; p)\, \hat{u}(x_1) + P_2(x; p)\, \hat{u}(x_2), \tag{1}$$

where $P_2(x; p) = 1 - P_1(x; p)$. Recall that $\hat{v}$ and v represent the same preferences if and only if each is a monotonic transformation of the other. Consider a strictly increasing (and for simplicity, differentiable) function $f : \mathbb{R} \to \mathbb{R}$, and define $\hat{v} = f \circ v$. Then, we seek a probability distribution $P(x; p)$ and a utility function over prizes $\hat{u}$ such that (1) is satisfied. By considering the case that the outcomes in the two states are the same (that is, the case of constant acts), note that (1) implies that $\hat{v}(z, z) = \hat{u}(z) = f(v(z, z)) = f(u(z))$ for all z. Then the desired representation (1) simplifies to

$$\hat{v}(x) = f(v(x)) = P_1(x; p) f(u(x_1)) + (1 - P_1(x; p)) f(u(x_2)).$$

Solving for $P_1(x; p)$, we get

$$P_1(x; p) = \frac{f(v(x)) - f(u(x_2))}{f(u(x_1)) - f(u(x_2))}$$

for $x_1 \neq x_2$. Note that $P_1(x; p)$ is always between zero and one because, by properties of expected utility, $v(x)$ is always between $u(x_1)$ and $u(x_2)$. As $x_1 \to x_2$, $P_1(x; p)$ converges to p_1. Naturally, $P_2(x; p) := 1 - P_1(x; p)$. When $x_1 > x_2$, the denominator of $P_1(x; p)$ is positive. Thus, when f is convex, Jensen's inequality implies that

$$P_1(x; p) \leq \frac{p_1 f(u(x_1)) + (1 - p_1) f(u(x_2)) - f(u(x_2))}{f(u(x_1)) - f(u(x_2))} = p_1.$$

The probability of the bigger prize is thus distorted down. Similarly, when f is concave, the

---

[4] Although Savage's original work applies only to the case where the state space is not finite, it has been shown how to derive a Savage-type representation when there are only a finite number of states (see, e.g., Wakker (1984) or Gul (1992)).

probability of the bigger prize is distorted up. (An analogous characterization holds when $x_2 > x_1$: the probability of the smaller prize is distorted up when f is convex, and distorted down when f is concave). Stated differently, the pessimist holds beliefs that are first-order stochastically dominated by the standard Savage distribution, while the optimist holds beliefs that first-order stochastically dominate it.

For specific classes of convex and concave functions, we can say more. Without loss of generality, we assume for the proposition below that the utility level $u(x)$ is positive for each $x \in X$.

Proposition 1. *Consider $x_1 \neq x_2$ and the transformation $f(z) = z^r$. Then $\frac{\partial P_i(x;p)}{\partial x_i} < 0$ for $r > 1$, and $\frac{\partial P_i(x;p)}{\partial x_i} > 0$ for $r \in (0,1)$.*[5]

The proof appears in the appendix. The case $r = 1$ corresponds to the standard Savage formulation in which there is no stake-dependent probability distortion. When $r > 1$, the DM's probability assessments reflect a stronger notion of pessimism. The better the consequence in any state, the less likely he thinks that this state will be realized. In particular, improving the best outcome reduces his assessment of its probability (as in the example in the introduction). Similarly, making the worst outcome even worse increases his assessment of its probability. When $r \in (0,1)$ the comparative statics are flipped. For the optimist, the better is the best outcome, the more likely the DM thinks it is; and the worse is the worst outcome, the less likely he thinks it is. By construction, however, choice behavior in either case is indistinguishable from that of a DM with a Savage-type representation.

Remark 1. *Based on Aumann's 1971 exchange of letters with Savage (reprinted in Drèze (1987)), the following argument has often been used to point out that the Savage representation could have multiple state-dependent expected utility representations, leaving the (single) probability distribution ill-defined. Consider a Savage representation of the form $p_1 u(x_1) + p_2 u(x_2)$. Then for any $\widehat{p}$ with the same support as p, this expression is equal to $\widehat{p}_1 \widehat{u}(x_1, s_1) + \widehat{p}_2 \widehat{u}(x_2, s_2)$, where $\widehat{u}(x_s, s) := \frac{p_s u(x_s)}{\widehat{p}_s}$ is a state-dependent utility function. This means that the Savage probability distribution is only unique under the assumption of state-independent utility. Notice that the same "multiply-and-divide" approach **cannot** be used to generate stake-dependent probabilities. To see this, fix a strictly positive utility function over prizes, $\overline{u}$ and let $\overline{p}(s, x_s) := \frac{p_s u(x_s)}{\overline{u}(x_s)}$. While the expected utility $p_1 u(x_1) + p_2 u(x_2)$ is the same as $\overline{p}(s_1, x_1)\overline{u}(x_1) + \overline{p}(s_2, x_2)\overline{u}(x_2)$, note that $\overline{p}(s, x_s)$ is not a probability distribution unless we normalize it by $\overline{p}(s_1, x_1) + \overline{p}(s_2, x_2)$. However, that scaling factor is not a constant – it generically depends on the stakes in all states (that is, unless u is a scalar multiple of $\overline{u}$). Therefore, the resulting utility representation no longer represents the same preferences as the original Savage representation.*

---

[5]One can find convex or concave functions outside this class for which the result does not hold. As an example, suppose $f(z) = 3z^2 - z^3$ if $z \in (0,1]$ and $f(z) = -1 + 3z$ for $z \in (1, \infty)$, which is a convex function. For $u(x) = x$, $p = (1/2, 1/2)$, $x_1 \in (0,1)$ and $x_2 = 1/4$, notice that $\frac{\partial P_1(x;p)}{\partial x_1} = \frac{1}{8} + \frac{27 + 18x_1}{44 + 16x_1(11 - 4x_1)} > 0$.

2.2 The general case

We have shown above how to construct a continuum of "expected utility" representations using distorted probabilities when there are two-states of nature. Under any of these representations, the certainty equivalent of each act is the same as that under the original Savage representation. While the computation of alternative representations is particularly simple in the two-state case, the multiplicity of representations does not depend on there being only two states. We next show this can be done for any finite number of states.[6]

Let $S = \{s_1, ..., s_n\}$ be the set of states and let $x = (x_1, ..., x_n) \in \mathbb{R}^n$ be an act, where x_i corresponds to the outcome in state s_i. Consider a Savage expected utility representation, with p the probability vector and u the utility function over prizes. We look for a stake-dependent probability distribution $P(x; p)$ and a representation of the form

$$\hat{v}(x) = \sum_{i=1}^{n} P_i(x; p) \hat{u}(x_i). \tag{2}$$

For $\hat{v}$ to represent the same preferences as the Savage expected utility function v, there must exist an increasing transformation f such that $\hat{v} = f \circ v$. As before, this implies that

$$\hat{v}(x) = f(v(x)) = \sum_{i=1}^{n} P_i(x; p) f(u(x_i)). \tag{3}$$

Including the above equation and the obvious restriction that $\sum_{i=1}^{n} P_i(x; p) = 1$, we have two equations with n unknowns. While this sufficed for a unique solution (given u and $\hat{u}$) in the case $n = 2$, when $n \geq 3$ there will generally be many ways to construct a probability distortion, corresponding to different ways a DM might allocate weight to events. More specifically, for an act x, let $ce(x; u, p)$ be the certainty equivalent of x given a utility function $u(\cdot)$ and a probability distribution p: $u(ce(x; u, p)) = \sum_{i=1}^{n} p_i u(x_i)$. Consider a transformation f which is convex (concave). Since $f \circ u$ is less risk averse than u, $ce(x; f \circ u, p) > ce(x; u, p)$ whenever x is non-degenerate (the reverse inequality holds if f is concave). We define

$$\mathcal{P}(x, p, u, f) = \left\{ q \in [0, 1]^n : \sum_{i=1}^{n} q_i = 1 \text{ and } ce(x; f \circ u, q) = ce(x; u, p) \right\}$$

to be the set of probability distributions with the property that for any $q \in \mathcal{P}(x, p, u, f)$, the certainty equivalent of $f \circ u$ with respect to the lottery q equals that of u with respect to the Savage distribution p. Thus $\mathcal{P}(x, p, u, f)$ is the set of probability distributions that for the given prizes yield expected utility equal to the certainty equivalent, that is, the indifference curve in the space of probabilities that corresponds to that expected utility. Figure 1a illustrates this with the Machina-Marschak triangle for the case $n = 3$ (the probability of the highest prize x_3 is on the vertical axis and the probability of the worst prize x_1 is on the horizontal axis). The line $\mathcal{P}(x, p, u, f)$ is

---

[6]Alternatively, it can be done for simple (finite support) acts on a continuum state space.

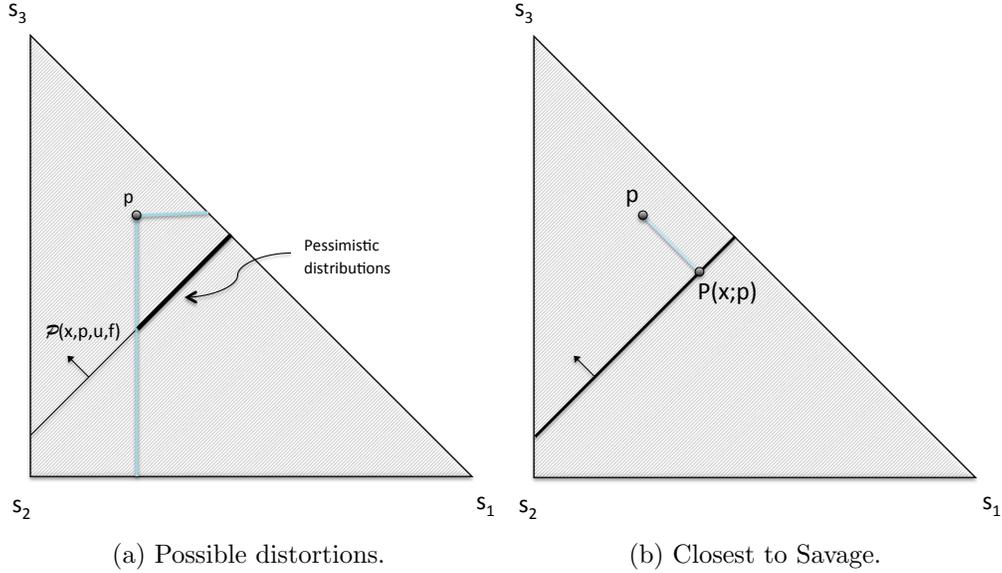(a) Possible distortions.　　(b) Closest to Savage.

Figure 1: Constructing a pessimistic representation in the case $n = 3$. In the figures, $x_3 > x_2 > x_1$. The x-axis represents p_1, the y-axis represents p_3, and $p_2 = 1 - p_1 - p_3$.

the set of probabilities for which expected utility is equal to $ce(x; u, p)$, and must pass through a point lying below p. Otherwise, the certainty equivalent of p under $f \circ u$ would be higher than the certainty equivalent under u. The distortions in the bolded portion of $\mathcal{P}(x, p, u, f)$ in Figure 1a are pessimistic: they lie southeast of p on the indifference curve $\mathcal{P}(x, p, u, f)$, and are thus both first-order stochastically dominated by p and deliver the same certainty equivalent under $f \circ u$ as does the Savage representation.

As is apparent from the figure, there are multiple ways to select a pessimistic probability distortion. We will demonstrate one simple mapping from acts to pessimistic beliefs. For any two probability distributions q, q' over S, let $d(q, q')$ be the Euclidean distance between them:

$$ d(q, q') = \sqrt{\sum_{i=1}^{n} (q_i - q_i')^2}. $$

We associate with any act the probability distribution in $\mathcal{P}(x, p, u, f)$ that is of minimal distance to the Savage distribution p:

$$ P(x; p) = \underset{q \in \mathcal{P}(x, p, u, f)}{\arg\min} \ d(p, q). \tag{4} $$

This mapping is illustrated in Figure 1b for the case $n = 3$ and convex f. Note that the Savage distribution p first-order stochastically dominates $P(x; p)$. This property is true for any convex f.[7] It can be analogously shown that for any concave f (the case of optimism), the probability

---

[7]This can be shown geometrically. Fix any act x and consider a Machina-Marschak triangle oriented as before. Given the orientation of the triangle, the slope of the line $\mathcal{P}(x, p, u, f)$ is positive. Consequently, if we draw both the vertical line from p down to $\mathcal{P}(x; p, u, f)$, as well as the horizontal line passing through p, the angle formed between

distribution $P(x;p)$ constructed according to (4) will first-order stochastically dominate the Savage distribution p. The argument above is valid independently of the ranking of the three prizes, that is, $P(x;p)$ is a continuous function of the act. Different rankings generate different indifference curves, but a pessimist will always shift weight (relative to Savage) towards bad states, and the optimist will always shift weight towards good states. The argument also holds for any finite number of states n.

3   Risky lotteries and uncertain acts

Consider a DM who is faced with choice objects, some of which are lotteries with objective probabilities and others that are subjective acts as in Savage. Suppose the DM's preferences are defined over the union of $\mathcal{L}^1$, the set of (purely objective) simple lotteries over the set of prizes X, and $\mathcal{F}$, the set of (purely subjective) Savage acts over X. On the subdomain of subjective acts, the DM satisfies the axioms of Savage, leading to a subjective expected utility representation with Bernoulli function v and probability distribution p. On the subdomain of objective lotteries, the DM satisfies the axioms of vNM, leading to an expected utility representation with Bernoulli function u. Each of the two subdomains contains deterministic outcomes: for any outcome $x \in X$, $\mathcal{F}$ contains the constant act that gives x in every state, and $\mathcal{L}^1$ contains the lottery that gives x with probability 1. It is natural to assume that the DM is indifferent between these two objects. More formally, suppose the DM's preferences $\succeq$ are defined over the domain $\Psi = \mathcal{L}^1 \cup \mathcal{F}$.[8] The assumptions above imply that for any $\xi \in \Psi$, there is an increasing transformation h such that the DM's preferences over the domain Ψ are represented by

$$U(\xi) = \begin{cases} h\left(\sum_x \pi(x) u(x)\right) & \text{for } \xi = \pi \in \mathcal{L}^1 \\ \sum_s p_s v(l(s)) & \text{for } \xi = l \in \mathcal{F} \end{cases}, \text{ where } h(u(x)) = v(x) \text{ for all } x \in X. \quad (5)$$

To discuss ambiguity aversion in the context of Equation (5), consider the hypothetical two-urn experiment introduced by Ellsberg (1961). There are two urns each containing 100 balls which could be black or red. The composition of Urn 1 (the ambiguous urn) is unknown. Urn 2 (the risky urn) contains exactly 50 red and 50 black balls. The DM can bet on the color of the ball drawn from an urn. Ellsberg predicts that given either urn, most people would be indifferent between betting on either red or black – indeed, by symmetry, it is reasonable to assume that the two

---

each of these lines and the indifference curve $\mathcal{P}(x; p, u, f)$ is less than $90°$. Thus, the point $P(x; p)$ for which the line between p and $P(x; p)$ is exactly $90°$ must lie to the southeast of p (if it exists; otherwise the closest point is on the boundary of the triangle, southeast of p).

[8]Note that this domain is essentially a strict subset of the domain of Anscombe and Aumann (1963), in which the outcome of an act in every state is an objective lottery. This domain is similar to the one used in Chew and Sagi (2008). Using their language, the sets $\mathcal{L}^1$ and $\mathcal{F}$ can be thought of as two different sources of uncertainty, on which the DM's preferences may differ. This domain allows us to talk about ambiguity while abstracting from the multistage feature of Anscombe and Aumann (1963)'s model.

colors are equally likely in Urn 1. Yet, he predicts that people would prefer bets based on Urn 2 to corresponding bets based on Urn 1, because they would prefer knowing the exact probability distribution. For a DM with a representation of the form (5), such a preference occurs if and only if v is more concave than u (as seen using Jensen's inequality).[9,10]

Note that it is impossible to keep u and v the same without assuming ambiguity neutrality, as preferences are entirely characterized by the utility for prizes under the conventions of the expected utility representation. A natural focal point, however, is for the DM's utility over prizes (which captures his tastes for the ultimate outcomes) to be consistent across the objective and subjective domains; that is, $u = v$. Simply put, the prizes are the same in both domains; it is only the probabilities that differ in the two situations.

Our model, on the other hand, may attribute the DM's ambiguity aversion (in the sense of preferring bets on the risky urn to bets on the ambiguous urn) to pessimistic probability assessments, rather than a change in his utility for prizes. Indeed, observe that if $U(\cdot)$ in Equation (5) represents the DM's preferences, so does $\widehat{U} = h^{-1} \circ U$. If $v = h \circ u$, where h is concave, we apply our previous results of Section 2 using the *convex* transformation h^{-1}. That is, there exists a pessimistic, stake-dependent probability distribution $P(\cdot; p)$ that is first-order stochastically dominated by the Savage distribution p, such that the utility representation $\widehat{U} = h^{-1} \circ U$ may be written as

$$\widehat{U}(\xi) = \begin{cases} \sum_x \pi(x) u(x) & \text{for } \xi = \pi \in \mathcal{L}^1 \\ \sum_s P_s(\xi; p) u(l(s)) & \text{for } \xi = l \in \mathcal{F} \end{cases}.$$

Analogously, if a DM is discovered to be "ambiguity loving," then he may be viewed as an optimist using the same utility function over prizes in both domains. Much in the same way that probabilities are identified under Savage's convention of *state*-independent utility for prizes (as discussed in Remark 1), we may detect the presence of optimism or pessimism under the convention of *source*-independent utility for prizes. In that sense, our approach is different than that of Chew and Sagi (2008), who also use source-dependent expected utility on a similar domain to address ambiguity aversion. In their work, ambiguity aversion is captured by the source-dependent curvature of the *utility* for prizes.[11]

---

[9]To see this, observe that by applying h^{-1}, $\widehat{U}(\xi) = \begin{cases} \sum_x \pi(x) u(x) & \text{for } \xi = \pi \in \mathcal{L}^1 \\ h^{-1}\left(\sum_s p_s h(u(l(s)))\right) & \text{for } \xi = l \in \mathcal{F} \end{cases}$, also represents the DM's preferences. For acts and lotteries with two possible outcomes as in Ellsberg's example, and for which $\pi(x) = p_s = \frac{1}{2}$ for every x, s, a direct application of Jensen's inequality says that $\sum_x \frac{1}{2} u(x) > h^{-1}\left(\sum_s \frac{1}{2} h(u(l(s)))\right)$ for all prizes $x, l(s)$ if and only if h^{-1} is convex, which in turn holds if and only if v is a concave transformation of u.

[10]This also corresponds to Ghirardato and Marinacci's (2002) definition of ambiguity aversion. According to their definition, the DM is more risk averse in uncertain settings than in objective settings if there exists a probability distribution p over S, such that for all $\pi \in \mathcal{L}^1$ and $l \in \mathcal{F}$, $l \succeq \pi$ implies that $\mu_{l,p} \succeq \pi$, where $\mu_{l,p}$ is the objective lottery under which the prize $l(s)$ is received with probability $p(s)$. The intuition behind this axiom is that if the DM prefers an act to a given lottery, it would also be better to simply receive that "act" with the objective probabilities that would ultimately be specified by the Savage distribution.

[11]The idea of capturing attitude towards ambiguity entirely through the utility function also appears in Klibanoff et al. (2005) and Ergin and Gul (2009), among others.

In principle, one may be able to use preferences on objective lotteries and subjective acts to determine whether a DM distorts probabilities, and suggest a comparative measure of pessimism. For example, take two decision makers with identical Savage preferences over subjective acts, that is, both individuals admit a Savage representation $\langle v, p \rangle$. However, they have different preferences over objective lotteries, with utility over prizes u_1 and u_2, respectively. Under the model above, each DM distorts probabilities differently, with DM_i's distortion function f_i given by $u_i \circ v^{-1}$. If f_1 is more convex than f_2 – or equivalently, u_1 is more convex than u_2 – then DM_1 is more pessimistic than DM_2. To illustrate using our minimum-distance example from Section 2.2 (see Equation (4)), define the level of pessimism of DM_i as $\min\limits_{q \in \mathcal{P}(x,p,v,f_i)} d(p,q)$, that is, the minimum distance between Savage's p and the point in the simplex that generates the same preferences given DM_i's distortion function f_i. Observe that u_1 is more convex than u_2 iff $\min\limits_{q \in \mathcal{P}(x,p,v,f_1)} d(p,q) > \min\limits_{q \in \mathcal{P}(x,p,v,f_2)} d(p,q)$ for any vector of prizes x. If $u_i = v$, that is, if the Savage and the vNM utility functions coincide, then DM_i is not pessimistic. This is in line with the idea that the Ellsberg paradox captures the relative, not absolute, extent of pessimism in the uncertainty domain versus the risk domain; see, for example, Wakker (2010, Chapter 11).

This argument, however, presumes that the DM takes the probabilities of objective lotteries at face value. Her choices will depend on the likelihoods in her mind of getting the various prizes. She may well think "I'm very unlikely to get the \$100 if I take the gamble - I never win anything." There is no compelling reason to believe that a pessimistic or optimistic DM's mental assessment of the likelihood of an event can be controlled by arguing what the DM "should" believe.

4 Discussion and related literature

4.1 Stake-dependent probabilities in other models

While our approach differs from that taken by other researchers, it is quite standard in the literature on ambiguity aversion to model the DM as though he evaluates outcomes according to expected utility, with an unvarying utility function and a probability distribution that depends on the outcome being evaluated. Consider, for example, one of the most widely known models of decision making under uncertainty, the maxmin expected utility with non-unique prior model of Gilboa and Schmeidler (1989). In their model, the DM behaves as though there is a set of possible probabilities that can be used, along with a fixed utility function, to compute the expected utility of any act. For any act, the probability used is the one that yields the lowest expected utility among those in their set. If the set of possible probabilities is a singleton, their model reduces to the standard model with stake-independent probabilities. A DM who is uncertain about the exact probability distribution to use (that is, a DM for whom the set of possible probabilities is not a singleton), will use probabilities that typically vary with the act in question. This is illustrated in Figure 2, where the shaded region is the set of probabilities the DM thinks possible. Orienting
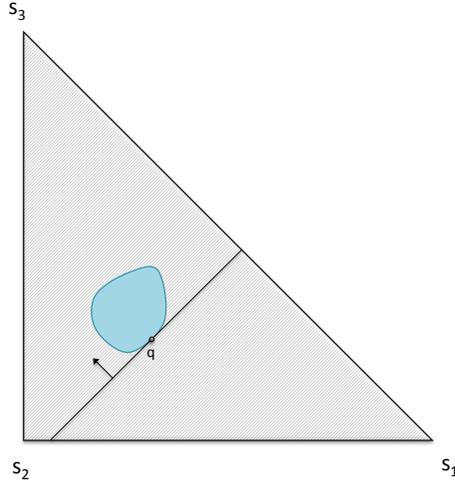
10

Figure 2: The indifference curve through the probability distribution that minimizes expected utility when $x_1 < x_2 < x_3$, as in Gilboa and Schmeidler (1989). If the indifference map becomes steeper (e.g., if x_3 decreases), then the distribution moves from q to a point further up the boundary of the probability set.

the Machina-Marschak triangle as before, with $x_3 > x_2 > x_1$, the probability that minimizes the expected utility over that set is q. If the prize x_3 decreases, the indifference map becomes steeper and the probability that minimizes expected utility over the same shaded set moves up along the boundary. Observe that when there are at least three states and the set of probabilities the DM thinks possible is strictly convex, there will be a continuous function that assigns to each act a unique, stake-dependent probability which the DM uses to compute expected utility, just as is the case with the "least distance" mapping described in the previous section.

Thus, both the maxmin expected utility model of Gilboa and Schmeidler (1989) and our model capture the choice behavior of agents who adapt the probability used in the expected utility calculation to the outcome being evaluated. There is, of course, a major difference between the two models: the choices of agents who employ the maxmin method of choosing probabilities will typically violate Savage's axioms on the subjective domain, while ours satisfy those axioms by construction. The consequence, of course, is that the maxmin expected utility model can generate behavior that cannot arise in our model.

4.2 Behavioral notions of optimism and pessimism

In this paper we discuss a cognitive notion of optimism and pessimism. A number of papers discuss optimism and pessimism as behavioral phenomena that are incompatible with expected utility. Wakker (1990), for example, defines pessimism through behavior (similarly to uncertainty aversion) and shows that within the rank-dependent expected utility (RDU) model, pessimism (optimism) holds if greater decision weights are given to worse (better) ranks. (See also Wakker (2001)). In con-

11

trast to our model, in Wakker's model changes in outcomes affect decision weights only when ranks change. Two recent papers also investigate behavioral notions of pessimism. Using the Anscombe and Aumann (1963) framework, Dean and Ortoleva (2012) suggest a generalized notion of hedging, which captures pessimism and applies to both objective risk and subjective uncertainty. Gumen, Ok, and Savochkin (2012) introduce a new domain which allows subjective evaluations of objective lotteries. They use their framework to define a general notion of pessimism for objective lotteries in a way reminiscent of uncertainty aversion for subjective acts. Their definition of pessimism is not linked to any specific functional form and hence applies to a broader class of preferences than just the RDU (as in Wakker). It also can incorporate stake-dependent probabilities.

4.3   Other related literature

The observation that the Savage-type representation and the optimist (or pessimist) can support the same underlying preferences, and hence cannot be distinguished by simple choice data, is related to general comments about model identification. In a series of papers, Karni (2011 and references therein) points out that the identification of probabilities in Savage's model rests on the (implicit) assumption of state-independent utility, and proceeds to propose a new analytical framework within which state independence of the utility function has choice-theoretic implications.[12] In the context of preference over menus of lotteries, Dekel and Lipman (2011) point out that a stochastic version of Gul and Pesendorfer (2001)'s temptation model is observationally equivalent to a random Strotz model. Chatterjee and Krishna (2009) show that a preference with a Gul and Pesendorfer (2001) representation also has a representation where there is a menu-dependent probability that the choice is made by the tempted (the "alter-ego") self, and otherwise the choice is made by the untempted self. Spiegler (2008) extends Brunnermeier and Parker's (2005) model of optimal expectations by adding a preliminary stage to the decision process, in which the DM chooses a signal from a set of feasible signals. Spiegler establishes that the DM's behavior throughout the two-stage decision problem, and particularly his choices between signals in the first stage, is indistinguishable from those of a standard DM who tries to maximize the expectation of some state-dependent utility function over actions. In the context of preferences over acts, Strzalecki (2011) shows that for the class of multiplier preferences, there is no way of disentangling risk aversion from concern about model misspecification. Consequently, he points out that "...policy recommendations based on such a model would depend on a somewhat arbitrary choice of the representation. Different representations of the same preferences could lead to different welfare assessments and policy choices, but such choices would not be based on observable data." Some of the papers above suggest additional choice data that is sufficient to distinguish between the models. We emphasize that our goal in this

---

[12]Grant and Karni (2005) argue that there are situations in which Savage's notion of subjective probabilities (which is based on the convention that the utilities of consequences are state-independent) is inadequate for the study of incentive contracts. For example, in a principal-agent framework, misconstrued probabilities and utilities may lead the principal to offer the agent a contract that is acceptable yet incentive incompatible.

paper is to show that, however one might interpret its canonical representation, the Savage model is consistent with a notion of pessimism. We do not aim at distinguishing our model from Savage's model.

Appendix

Proof of Proposition 1

We start with some mathematical preliminaries. Consider $\{s_1, s_2\}$. Suppose that s_1 occurs with probability p_1 and s_2 occurs with probability $p_2 = 1 - p_1$. Let $a \neq b$ be two positive real numbers. Define two random variables, X and Y as follows: X has value a in state s_1 and b in state s_2; and Y has value a in state s_2 and b in state s_1. We claim that for any number s,

$$(ab)^s = E\left[X^{-s}\right]^{-1} E\left[Y^s\right]. \tag{6}$$

To show this, note that

$$
\begin{aligned}
E\left[X^{-s}\right]^{-1} E\left[Y^s\right] &= \left(p_1 a^{-s} + p_2 b^{-s}\right)^{-1} \left(p_1 b^s + p_2 a^s\right) \\
&= \frac{\left(p_1 b^s + p_2 a^s\right)}{\left(p_1 a^{-s} + p_2 b^{-s}\right)} = \frac{a^s b^s \left(p_1 b^s + p_2 a^s\right)}{\left(p_1 b^s + p_2 a^s\right)} = (ab)^s.
\end{aligned}
$$

We focus on the derivative of $P_1(x_1, x_2; p_1)$ with respect to x_1, since the other case is identical. Taking the derivative and simplifying, we find that using the transformation $f(z) = z^r$, equals

$$\frac{\partial P_1(x_1, x_2; p_1)}{\partial x_1} = \frac{r u'(x_1) \left[u(x_1)^r u(x_2)^r - (p_1 u(x_1) + p_2 u(x_2))^{r-1} (p_2 u(x_1)^r u(x_2) + p_1 u(x_2)^r u(x_1))\right]}{u(x_1)(u(x_1)^r - u(x_2)^r)^2}.$$

Since $r, u'(x_1), u(x_1), u(x_2) > 0$, the sign of $\frac{\partial P_1(x_1, x_2; p_1)}{\partial x_1}$ equals the sign of

$$u(x_1)^r u(x_2)^r - (p_1 u(x_1) + p_2 u(x_2))^{r-1} (p_2 u(x_1)^r u(x_2) + p_1 u(x_2)^r u(x_1)).$$

Let $a = u(x_1)$ and $b = u(x_2)$. Factoring out ab, the last expression has the sign of

$$a^{r-1} b^{r-1} - E\left[X\right]^{r-1} E\left[Y^{r-1}\right].$$

Using (6) with $s = r - 1$, this is equivalent to $E\left[X^{1-r}\right]^{-1} E\left[Y^{r-1}\right] - E\left[X\right]^{r-1} E\left[Y^{r-1}\right]$, which has the same sign as $E\left[X^{1-r}\right]^{-1} - E\left[X\right]^{r-1}$.

For the case $r > 1$, we would like to show that $E\left[X^{1-r}\right]^{-1} - E\left[X\right]^{r-1} < 0$, or equivalently, that $E\left[X^{1-r}\right]^{-1} < E\left[X\right]^{r-1}$. Applying Jensen's inequality to the convex transformation $g(x) = x^{1-r}$, we get $E\left[X^{1-r}\right] > E\left[X\right]^{1-r}$, or $E\left[X^{1-r}\right]^{-1} < E\left[X\right]^{r-1}$. For the case $r \in (0, 1)$, we want to show that $E\left[X^{1-r}\right]^{-1} - E\left[X\right]^{r-1} > 0$, or equivalently, that $E\left[X^{1-r}\right]^{-1} > E\left[X\right]^{r-1}$. Applying Jensen's inequality to the concave transformation $g(x) = x^{1-r}$, we get $E\left[X^{1-r}\right] < E\left[X\right]^{1-r}$, or $E\left[X^{1-r}\right]^{-1} > E\left[X\right]^{r-1}$. ∎

References

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

Ariely, Dan (2010), "The Upside of Irrationality: The Unexpected Benefits of Defying Logic at Work and at Home," Harper (publisher).

Brunnermeier, Markus K. and Jonathan Parker (2005), "Optimal Expectations," *American Economic Review,* 95(4): 1092-1118.

Chatterjee, Kalyan and R. Vijay Krishna (2009), " A 'Dual Self' Representation for Stochastic Temptation," *American Economic Journal: Microeconomics,* 1: 148-167.

Chew, Soo Hong and Jacob Sagi (2008), "Small Worlds: Modeling Attitudes Toward Sources of Uncertainty," *Journal of Economic Theory,* 139: 1-24.

Dean, Mark and Pietro Ortoleva (2012), "Allais, Ellsberg, and Preferences for Hedging," *Mimeo, California Institute of Technology.*

Dekel, Eddie and Bart Lipman (2010), "How (Not) to Do Decision Theory," *Annual Review of Economics* 2: 257-82.

Dekel, Eddie and Bart Lipman (2011), "Costly Self Control and Random Self Indulgence", mimeo.

Drèze, Jacques H. (1987), "Decision Theory with Moral Hazard and State-Dependent Preferences," in *Drèze, J. H. Essays on Economic Decisions Under Uncertainty.* Cambridge University Press, Cambridge.

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

Ergin, Haluk and Faruk Gul (2009), "A Theory of Subjective Compound Lotteries," *Journal of Economic Theory,* 144(3): 899-929.

Gilboa, Itzhak and David Schmeidler (1989), "Maxmin expected utility with non-unique prior," *Journal of Mathematical Economics,* 18(2): 141-153.

Grant, Simon and Edi Karni (2005), "Why Does It Matter that Beliefs and Valuations be Correctly Represented?" *International Economic Review* 46: 917-934.

Grant, Simon, Ben Polak and Tomasz Strzalecki (2009), "Second-Order Expected Utility," *mimeo.*

Gul, Faruk (1992), "Savage's Theorem With a Finite Number of States," *Journal of Economic Theory 57(1).*

Gul, Faruk and Wolfgang Pesendorfer (2001), "Temptation and Self-Control," *Econometrica* 69: 1403–1436.

Gumen, Anna, Efe Ok and Andrei Savochkin (2012), "Decision-Making under Subjective Risk: Toward a General Theory of Pessimism," *Mimeo, New York University.*

Hey, John (1984), "The Economics of Optimism and Pessimism," *Kyklos* 37: 181-205.

Karni, Edi (2009), "A Theory of Medical Decision Making Under Uncertainty," *Journal of Risk and Uncertainty* 39: 1-16.

Karni, Edi (2011), "Subjective Probabilities on a State Space," *American Economic Journal: Microeconomics*, forthcoming.

Klibanoff, Peter, Massimo Marinacci and Sujoy Mukerji (2005), "A Smooth Model of Decision Making Under Ambiguity " *Econometrica* 73(6): 1849-1892.

Pesendorfer, Wolfgang (2006), "Behavioral Economics Comes of Age: A Review Essay on Advances in Behavioral Economics," *Journal of Economic Literature* 44: 712-721.

Savage, L. J. (1954), *The Foundation of Statistics*, Wiley, New York.

Spiegler Rani (2008), "On Two Points of View Regarding Revealed Preferences and Behavioral Economics," *The Foundations of Positive and Normative Economics*, OUP: 95-115.

Strzalecki, Tomasz (2011), "Axiomatic Foundations of Multiplier Preferences," *Econometrica* 79: 47-73.

von Neumann, John (1955), 'Method in the Physical Sciences', *Collected Works*, vol. 6, 491-498; originally in L. Leary (Ed.): *The Unity of Knowledge*, New York: Doubleday, 157-164.

Wakker, Peter (1984), "Cardinal Coordinate Independence for Expected Utility," *Journal of Mathematical Psychology* 28: 110-117.

Wakker, Peter (1990), "Characterizing Optimism and Pessimism Directly Through Comonotonicity," *Journal of Economic Theory* 52: 453-463.

Wakker, Peter (2001), "Testing and Characterizing Properties of Nonadditive Measures Through Violations of the Sure-Thing Principle," *Econometrica* 69: 1039-1059.

Wakker, Peter (2008), "Lessons Learned by (from?) an Economist Working in Medical Decision Making," *Medical Decision Making* 28: 690-698.

Wakker, Peter (2010), "Prospect Theory For Risk and Uncertainty," Cambridge University Press, Cambridge, UK.