Optimism and Pessimism with Expected Utility*

David Dillenberger
University of Pennsylvania

Andrew Postlewaite†
University of Pennsylvania

Kareen Rozen‡
Brown University

March 2016

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. 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. Extending the DM's choice domain to include both subjective acts and objective lotteries, we consider a DM who behaves in accordance with expected utility on each subdomain, applies the same Bernoulli utility function over prizes regardless of their source, but may be optimistic or pessimistic with regard to subjective acts. This model can accommodate, for instance, the behavior in Ellsberg's two-urn experiment, and provides a framework within which optimism, pessimism, and standard Savage agents can be distinguished.

*First version August 2011. We thank Eddie Dekel, Itzhak Gilboa, Edi Karni, Mark Machina, Larry Samuelson, Tomasz Strzalecki, and Peter Wakker for helpful discussions and suggestions.

†Postlewaite thanks the National Science Foundation for support.

‡Rozen thanks the National Science Foundation for support under grant SES-0919955.

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, the DM 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. The DM 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, the 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, the 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 $p(y_1, y_2)$. The 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 distribution, 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 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.

Would a richer choice data allow us to distinguish optimistic, pessimistic, and standard Savage agents? We address this question by enriching the DM's choice domain to contain both objective lotteries and subjective acts. We propose an extended model in which the DM behaves in accordance with expected utility on each subdomain, applies the same Bernoulli utility function u over final prizes regardless of their source, but may be optimistic or pessimistic with regard to subjective acts. This model can, for instance, 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.[1] We show that optimism and pessimism can be identified by the extent to which the Bernoulli u from the objective domain is more or less convex than Savage's imputed utility v. As in classical theory, u encapsulates the DM's risk aversion on the *objective* domain, where it is applied to fixed, stake-independent probabilities. By contrast, the DM's attitude to gambles in the subjective domain is the net effect of two forces: the curvature of u and the stake-dependent distortion of probabilities.

It is valuable to know that preferences that satisfy Savage's axioms have an alternative representation to the standard expected utility with subjective probability beliefs. Economists have a set of models that they use to structure how they look at economic problems. When behavior that is at odds with those models comes to mind the inclination is to modify current models or lay out new models that accommodate the new behavior. An economist hearing of an agent who describes his decision-making process as in the example above would likely think that it was necessary to formulate a new model if she wanted accommodate this agent's decisions.[2] It is important to understand that there are many representations that satisfy Savage's axioms. Rather than setting out a new model to accommodate the behavior of the pessimistic agent above, the economist should understand that the behavior is consistent with the standard Savage model and that the range of economics that assumes standard subjective expected utility applies to this agent as well. This is not to say that *all* pessimistic behavior is consistent, only that some are. If one wished to model more precisely a particular form of pessimism, one might well need to go outside the subjective expected utility framework.

It is thus 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

[1] 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)).

[2] 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.

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 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.[3,4]

Consider our model in a situation where the 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 hypothetical choices 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.

[3] 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.

[4] 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.[5] 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

[5] 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)). Axioms for such a representation are presented in the Appendix.

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 $r \in (0,1)$.*[6]

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.

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.[7]

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

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

[6] 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 $\frac{\partial P_1(x;p)}{\partial x_1} = \frac{1}{8} + \frac{27+18x_1}{44+16x_1(11-4x_1)} > 0$.

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

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 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:

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

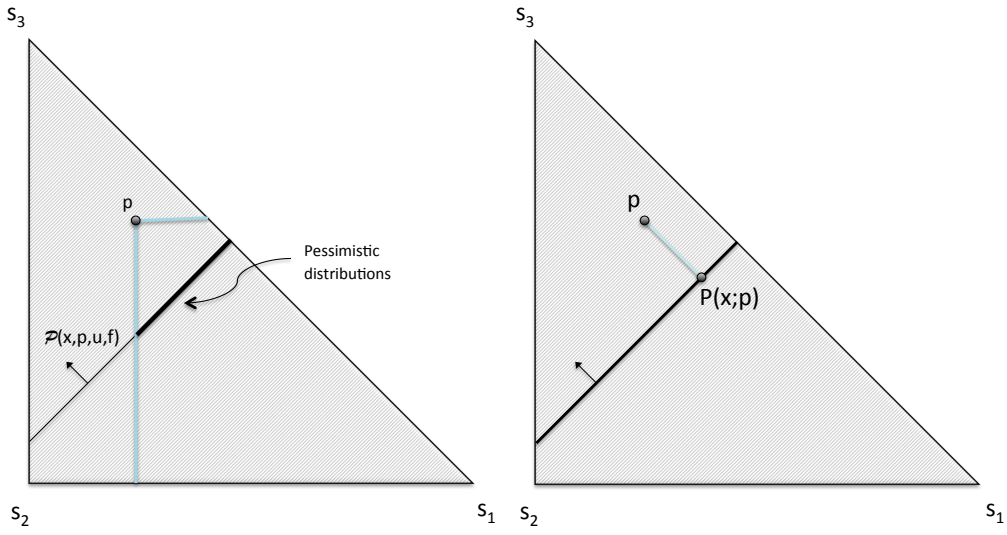

(a) Possible distortions. (b) Closest to Savage.

Figure 1: Constructing a pessimistic representation in the case $n = 3$. The horizontal axis represents the probability p_1 of state s_1, and the vertical axis represents the probability p_3 of state s_3. The probability of state s_2 is then $p_2 = 1 - p_1 - p_3$. In the figures, the stakes satisfy $x_3 > x_2 > x_1$, so that indifference curves to the northwest correspond to higher utility, as indicated by the arrow.

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.[8] It can be analogously shown that for any concave f (the case of optimism), the probability distribution $P(x; p)$ constructed according to (4) will first-order stochastically dominate the Savage distribution p. The construction 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 Identifying optimism and pessimism

In the previous section, we showed how behavior consistent with the model of expected utility is also consistent with our notions of optimism and pessimism. Two questions arise. First, is it possible to distinguish an optimist from a pessimist, and from a standard Savage agent? Second,

[8]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 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).

how do optimism and pessimism differ from risk aversion? In this section, we answer these questions by allowing the DM's choice domain to contain both objective lotteries and subjective acts. We propose an extended model in which the DM behaves in accordance with expected utility on each subdomain, applies the same Bernoulli utility function u over final prizes regardless of their source, but may be optimistic or pessimistic with regard to subjective acts.

Formally, the DM's complete and transitive preference relation $\succeq$ is defined over the union of $\mathcal{L}$, 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.[9] Each of these two subdomains contains deterministic outcomes: for any outcome $x \in X$, $\mathcal{F}$ contains the constant act that gives x in every state, and $\mathcal{L}$ contains the lottery that gives x with probability 1. It is natural to assume that the DM is indifferent between these two. Our DM also satisfies the axioms of Savage on the subdomain $\mathcal{F}$ of subjective acts, and the axioms of von Neumann-Morganstern on the subdomain $\mathcal{L}$ of objective lotteries. As a first step towards our model, note that these assumptions are immediately equivalent to the existence of two Bernoulli utility functions $u : X \to \mathbb{R}$ and $v : X \to \mathbb{R}$, an increasing transformation function $h : \mathbb{R} \to \mathbb{R}$, and a Savage probability distribution p such that the DM's behavior can be represented as maximizing the utility function $V : \mathcal{L} \cup \mathcal{F} \to \mathbb{R}$ given by:

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

Let us think about Equation (5) in the context of the two-urn thought 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 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. As seen using Jensen's inequality, such a preference occurs if and only if v is more concave than u.[10] Equivalently, since $h \circ u = v$, the transformation h must be concave. If, on the other hand, the DM prefers betting on Urn 1, then v is more convex than u, and the transformation h must be convex.

[9] 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 $\mathcal{L}$ 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.

[10] To see this, note that $h^{-1}(U(\xi)) = \widehat{U}(\xi) = \begin{cases} \sum_x \pi(x)\, u(x) & \text{for } \xi = \pi \in \mathcal{L}^1 \\ h^{-1}\left(\sum_s p_s h(u\left(l(s)\right))\right) & \text{for } \xi = l \in \mathcal{F} \end{cases}$ 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 $\sum_x \frac{1}{2} u(x) > h^{-1}\left(\sum_s \frac{1}{2} h(u\left(l(s)\right))\right)$ for all prizes $x, l(s)$ if and only if h^{-1} is convex, or v is a concave transformation of u.

Thus, in the context of Equation (5), modelling a DM who isn't neutral to the source of his gamble requires assuming that u and v differ.[11] A natural focal point, however, is for his utility over prizes (which capture his tastes for the ultimate outcomes) to be consistent across the objective and subjective domains.[12] After all, the prizes are the same across both domains; it is only the probabilities that differ. Our results from Section 2 make it possible to capture these features. We can model the DM as having the same utility for prizes regardless of the source, but simply being optimistic or pessimistic in his probability assessments. If we observe Ellsberg's predicted behavior, then the transformation h is concave, and our previous results from Section 2 identify the DM as a pessimist. This can be seen 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 $U = h^{-1} \circ V$ may be written as

$$U(\xi) = \begin{cases} \sum_x \pi(x)\, u(x) & \text{for } \xi = \pi \in \mathcal{L} \\ \sum_s P_s(\xi; p)\, u(l(s)) & \text{for } \xi = l \in \mathcal{F} \end{cases}. \tag{6}$$

Analogously, if the DM's choices reveal that h is convex, then the results of Section 2 identify him as an optimist. In that case, the stake-dependent probability distribution $P(\cdot; p)$ in Equation (6) would first-order stochastically dominate the Savage distribution p.

Thus, optimism and pessimism are not identified by the curvature of the classical Bernoulli utility u, but rather by the extent to which u is more or less convex than Savage's imputed utility v. As in classical theory, u encapsulates the DM's risk aversion on the *objective* domain, where it is applied to fixed, stake-independent probabilities. By contrast, the DM's attitude to gambles in the subjective domain is the net effect of two forces: the curvature of u and the stake-dependent distortion of probabilities. One could thus imagine a DM who feels comfortable with roulette wheels but less so with horse races. This closely relates to Ghirardato and Marinacci's (2002) definition of ambiguity aversion, which compares the DM's behavior in certain versus uncertain settings. The intuition behind their definition is that if the DM prefers an act to a given lottery, it would also be better to simply receive an 'objective version' of that act, in which the objective probabilities are those specified by the Savage distribution.[13] In situations where the DM's behavior still conforms with the axioms of Savage, our model identifies the underlying source of ambiguity attitude as optimistic or pessimistic probability distortion.

[11] Andreoni et al. (2015) describe experiments in a framework that is similar to ours in which subjects are faced with choices involving subjective and objective mixtures. Subjects exhibit inconsistencies that the authors suggest indicate a "directed pessimism," in which subjects are substantially more pessimistic about a subjective bet when it is mixed with higher sure outcomes than low outcomes.

[12] In that sense, our approach is different than that of Chew and Sagi (2008), who 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. 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.

[13] Formally, they say that 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 $\pi \in \mathcal{L}$ 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 $p(s)$.

Finally, note that one may be able to use preferences on objective lotteries and subjective acts not only to determine whether a DM distorts probabilities, but also to suggest a comparative measure of optimism or pessimism. For example, take two people with identical Savage preferences over subjective acts, that is, both 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 our model, 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)), we can define the level of pessimism of DM_i as $\min_{q \in \mathcal{P}(x,p,v,f_i)} d(p,q)$, 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

$$\min_{q \in \mathcal{P}(x,p,v,f_1)} d(p,q) > \min_{q \in \mathcal{P}(x,p,v,f_2)} d(p,q)$$

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 a standard Savage agent.

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 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

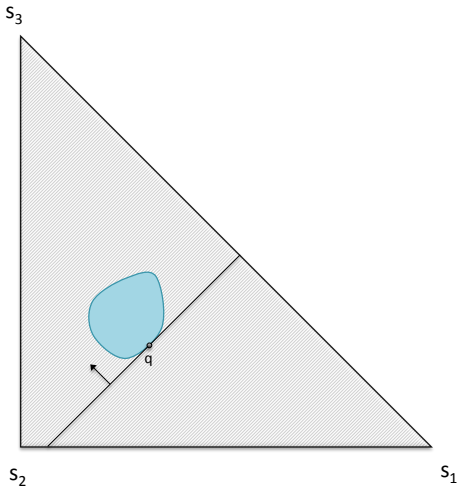

Figure 2: The indifference curve through the probability distribution that minimizes expected utility when $x_1 < x_2 < x_3$, as in Gilboa and Schmeidler (1989). If the indifference map becomes steeper (e.g., if x_3 decreases), then the distribution moves from q to a point further up the boundary of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

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 maximin 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 contrast 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 optimistic (or pessimistic) representation can describe the same underlying preferences, and hence cannot be distinguished by simple choice data, is related to general comments about model identification. Aumann's 1971 exchange of letters with Savage, reprinted in Drèze (1987), points out that the identification of probabilities in Savage's model rests on the implicit assumption of state-independent utility.[14] In a series of papers, Karni (2011 and references therein) addresses this issue by proposing a new analytical framework within which state independence of the utility function has choice-theoretic implications.[15] 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, '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 in question. Our goal in this pa-

[14] Notice that Aumann's multiply-and-divide approach for generating state-dependent representations (i.e., $p_1u(x_1) + p_2u(x_2) = \hat{p}_1 \frac{p_1u(x_1)}{\hat{p}_1} + \hat{p}_2 \frac{p_2u(x_2)}{\hat{p}_2} = \hat{p}_1\hat{u}_1(x_1) + \hat{p}_2\hat{u}_2(x_2)$) cannot generate stake-dependent probabilities. Normalizing the constructed factors to be probabilities will require dividing by something that generically depends on the stakes involved, so that the resulting utility representation no longer represents the same preferences as the original Savage representation.

[15] 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.

per is to show that, however one might interpret its canonical representation, the Savage model is consistent with notions of optimism and pessimism, and that predictions made within Savage's framework apply to such decision makers as well. As seen in Section 3, it becomes possible to identify optimism, pessimism and standard Savage decision makers once one extends the choice domain to include objective lotteries.

Appendix

A 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,

$$(ab)^s = E\left[X^{-s}\right]^{-1} E\left[Y^s\right]. \tag{7}$$

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

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

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

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

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

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

Using (7) with $s = r - 1$, this is equivalent to $E\left[X^{1-r}\right]^{-1} E\left[Y^{r-1}\right] - E[X]^{r-1} E\left[Y^{r-1}\right]$, which has the same sign as $E\left[X^{1-r}\right]^{-1} - E[X]^{r-1}$.

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

B Underlying Axioms for SEU

In this Appendix we outline the axioms underlying Savage's subjective expected utility representation. Since we assume a finite state space and a continuum of possible prizes, we will follow the approach of Gul (1992) (See Footnote 5)

Let Ω be a finite state space, $X = [w, b] \subset \mathbb{R}$ an interval of monetary prizes, and F be the set of all savage acts (that is, mappings from Ω to X). We identify with x the constant act that yields the same prize x in each state. For any event $A \subset \Omega$ and prizes $x, y \in X$, let xAy be the act defined as

$$[xAy](s) = \begin{cases} x & \text{if } s \in A \\ y & \text{if } s \notin A \end{cases}$$

A binary relation $\succeq$ is defined on F. The symmetric and asymmetric parts of $\succeq$ are denoted by $\succ$ and $\sim$, respectively. On $\succeq$ we assume the following axioms (for motivation of the axioms, we refer the reader to Gul (1992)).

Axiom 1. *$\succeq$ is complete and transitive.*

Axiom 2. *$f'(s) \sim f(s)Ah(s)$, $g'(s) \sim g(s)Ah(s)$ for all $s \in \Omega$ and A is not null*[16] *implies*

$$f \succ g \Leftrightarrow f' \succ g'$$

Axiom 3. *$x > y$ implies $x \succ y$. Furthermore, there exists $A^* \subset \Omega$ such that $xA^*y \sim yA^*x$ for all $x, y \in X$.*

Axiom 4. *For all $f \in F$, the sets $B(f) = \{g \in F : g \succeq f\}$ and $W(f) = \{g \in F : f \succeq g\}$ are closed.*[17]

Theorem 1 (Gul, 1992)**.** *If $\succeq$ satisfies Axioms 1-4, then there exists a probability measure p on the set of subsets of Ω and a function $u : X \to \mathbb{R}$ such that*

1. *$f \succeq g \iff \sum_s p(s)u(f(s)) \geqslant \sum_s p(s)u(g(s))$;*
2. *u is continuous and strictly increasing;*
3. *if (1) above holds when p is replaced by the probability measure p' and u is replaced by $u' : X \to \mathbb{R}$, then $p = p'$ and $u' = au + b$ for some $a > 0, b \in \mathbb{R}$.*

[16] An event A is null if $f(s) = g(s)$ for all $s \notin A$ implies $f \sim g$.

[17] Since $X \subset \mathbb{R}$, we can view F as a subset of $\mathbb{R}^{|\Omega|}$. A subset $G \subset F$ is closed if it is a closed subset of $\mathbb{R}^{|\Omega|}$.

References

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

Andreoni, James, Tobias Schmidt and Charles Sprenger (2015), "Measuring Ambiguity Aversion: Experimental Tests of Subjective Expected Utility," *Mimeo, UCSD.*

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 (2012), "Costly Self Control and Random Self Indulgence", *Econometrica* 80: 1271-1302.

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*, 3: 172-185.

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.