

Ambiguity in Game Theory?

Ken Binmore
Economics Department
Bristol University
8 Woodland Road
Bristol BS8 1TB, UK


Abstract: Various modifications of rational choice theory have sought to accommodate reported deviations from Bayesian decision theory in experiments on the Ellsberg paradox. The new theories usually postulate some ambiguity in the probabilities assigned to uncertain events. How well do such theories work when applied in game theory?

This question is explored from the viewpoint of Leonard Savage, who argued that his newly created theory of subjective expected utility is only realistically applicable in what he called a small world. The world of classical game theory is small almost by definition, but we endorse the view that the classical definition of a mixed strategy needs to be expanded to allow for devices from the theory of algorithmic randomization that allow pure strategies to be muddled together in a manner that defies a Bayesian description. However, we depart from the developing orthodoxy on game theory in the presence of Knightian uncertainty in also denying that an ambiguity description in terms of multiple priors is always adequate. Even when such an ambiguity description is adequate, it is argued that only an ambiguity neutral version of the Hurwicz criterion is viable, and so the common use of the maximin criterion in evaluating the payoffs that result when the players use Knightian strategies modeled in terms of multiple priors is excluded.

The paper continues by offering an argument that favors replacing the standard additive Hurwicz criterion by a multiplicative Hurwicz alternative in a game theory context. An example shows that both the ambiguity-neutral additive Hurwicz criterion and its multiplicative counterpart can generate muddled Nash equilibria that are Pareto-improvements on classical mixed Nash equilibria.

Key Words: Bayesian decision theory. Upper and lower probability. Hurwicz criterion. Alpha-maximin. Ellsberg game. Ellsberg strategy.

JEL Classification: C72, D81.



Inequality of the Arithmetic and Geometric Means

Ambiguity in Game Theory?¹

by Ken Binmore

1 Preview

Bayesianism is the ruling paradigm for rational choice behavior in risky or uncertain situations. The theory was created in 1954 by Leonard Savage [67] in his ground-breaking *Foundations of Statistics*, but sixty years later his warning on page 16 that it would be “preposterous” and “utterly ridiculous” to use his theory outside what he later calls a “small world” continues to be largely ignored. Nevertheless, experiments on the Ellsberg paradox have fuelled a literature that seeks to address the problems that had already been raised by Frank Knight [52] in 1921 about the consequences of proceeding as though one can always assign a precise probability to any uncertain event.²

Rational game theory stands clear of this issue largely because a game as classically conceived is a small world in Savage’s sense almost by definition. But should we not be trying to go beyond this classical vision to create a game theory that is adequate for applications in macroeconomics or finance? Is the current orthodoxy even adequate for the purposes for which it was invented in the first place? I share the view of a number of recent authors who take the radical line on both questions.³ However, this paper differs from most of this literature in holding that an ambiguity approach to the problem of Knightian uncertainty is not always adequate. Sometimes the problem is not that there is an ambiguity about what probability to attach to an event, but that it is not meaningful to attach a probability to the event at all—a fact that can occasionally be exploited for strategic purposes.

In applying ambiguity theories to game theory, it is usual to model the players as being ambiguity averse. The standard approach is to follow Wald [79] or

¹I am grateful for funding from the European Research Council under the European Community’s Seventh Framework Programme (FP7/2007-2013)/ERC grant 295449.

²For example, Berger [7], Breese and Fertig [13], Chrisman [18], DeGroot [20], Dempster [21], Earman [23], Fertig and Breese [28], Fine [29, 30], Fishburn [31, 32], Gardenfors and Sahlin [33], Ghirardato and Marinacci [35], Gilboa [36], Giron and Rios [38], Good [39], Halpern and Fagin [41], Huber [45], Kyburg [54], Levi [56, 57], Klibanoff *et al* [51], Pearl [64], Seidenfeld [69], Shafer [70], Wakker [76], Walley and Fine [81], and Walley [80].

³For example, Azrieli and Teper [4], Bade [6], Eichberger and Kelsey [24], Marinacci [61], Riedel and Sass [66].

Gilboa and Schmeidler [33] in using the maximin criterion to resolve ambiguities. Hurwicz [46], Chernoff [17], Milnor [62], Arrow and Hurwicz [2], and Ghirardato *et al* [34] offer a variety of reasons for generalizing the maximin criterion to the Hurwicz criterion when evaluating a simple gamble (in which the decision-maker wins when an event E occurs and loses otherwise). When it is known that E cannot have a probability less than a lower probability $\underline{p}(E)$ or more than an upper probability $\bar{p}(E)$, Hurwicz proposed the maximization of a weighted arithmetic mean of the upper and lower probabilities:

$$\pi_A(E) = (1 - \alpha)\underline{p}(E) + \alpha\bar{p}(E) \quad (1)$$

which reduces to the maximin criterion when $\alpha = 0$. However, with the ambiguity approach to Knightian uncertainty, we argue that the Hurwicz criterion is viable only in the ambiguity-neutral case when $\alpha = \frac{1}{2}$ (Theorem 1).

The paper continues by offering an approach to independence in the presence of Knightian uncertainty that suggests replacing the weighted arithmetic mean $\pi_A(E)$ by the corresponding geometric mean

$$\pi_G(E) = \{\underline{p}(E)\}^{1-\alpha} \{\bar{p}(E)\}^\alpha. \quad (2)$$

Neither version of the Hurwicz criterion retains the additive property of the probabilities of disjoint events, but the arithmetic version holds on to more of the flavor of the additive property than the geometric version. However, the geometric version retains the multiplicative property of the probabilities of independent events, which can be especially useful in game theory applications.

A simple 2×2 game with a unique Nash equilibrium in the classical sense is used to illustrate the relevance that these notions can sometimes have in game theory. In this game, the players get no more than their security payoffs at the equilibrium. Harsanyi [42, 43] and Aumann and Maschler [3] have independently asked why the players do not deviate to their security strategies in such games so as to get their equilibrium payoffs for sure. Enlarging the small world of mixed strategies by introducing “muddled strategies” provides a possible solution to this problem. In a muddled strategy, a player’s pure strategies are played with differing upper and lower probabilities. If outcomes of the game that result from the use of such muddled strategies are evaluated using either the arithmetic or the geometric Hurwicz criterion with $\alpha = \frac{1}{2}$, it is shown that non-classical Nash equilibria exist that Pareto-dominate the classical equilibrium.

Such muddled strategies are not implementable using an Ellsberg experimental set-up because their properties are inconsistent with their being assigned any

particular probability between their upper and lower probabilities. However, their introduction raises troublesome issues in the theory of algorithmic randomization.

2 Savage's Small Worlds

Why does Savage [67] restrict Bayesian decision theory to small worlds? The reason is that his crucial axioms demand that decisions be made *consistently* with each other. Nobody denies that a perfectly rational person would make decisions consistently, but Savage's theory is about subjective probabilities rather than what philosophers call epistemic probabilities that measure rational degrees of belief. If an adequate theory of this kind were available, consistency would come as part of the package, but no such solution of the problem of scientific induction is even remotely in sight.⁴

The Von Neumann and Morgenstern theory operates with objective probabilities, which we identify with limiting frequencies. Savage extended the Von Neumann and Morgenstern theory to cases when objective probabilities are unavailable. Their place is taken by his subjective probabilities that may differ from person to person because they simply reflect what bets a consistent decision-maker is willing to make.

Savage's theory therefore merely describes a person's choice behavior if that person makes choices consistently. But why would a person behave consistently? After all, physicists know that quantum theory and relativity are inconsistent where they overlap, but they live with this inconsistency rather than abandon the accurate predictions that each theory provides within its own domain. Physicists strive to create a "theory of everything" from which such inconsistencies have been removed, but everybody recognizes that this is a problem of enormous difficulty. So why do naive Bayesians proceed as though rationality somehow endows us with the talent to make consistent decisions without any effort at all?

Savage [67, p. 16] confined the field of application of his theory to small worlds because he did not believe that consistency comes any easier to us than it does to physicists. Calling Alice rational does not somehow endow her with a consistent set of beliefs (which Savage's theory then shows can be summarized with a prior distribution). Achieving consistency is part of what rationality is about—perhaps the most important part. We are entitled to appeal to Savage's

⁴If a solution to the problem of scientific induction ever comes in sight, my guess is that branches of mathematics much more complicated than probability theory will be required.

axioms only when this part of a rational analysis is over.

Savage's famous encounter with Maurice Allais in Paris illustrates how he thought Bayesian decision theory should be used in practice. When it was pointed out that his answers to what is now called the Allais paradox were inconsistent, he revised them until they were consistent. Luce and Raiffa [60, p. 302] summarize Savage's views as follows:

Once confronted with inconsistencies, one should, so the argument goes, modify one's initial decisions so as to be consistent. Let us assume that this jockeying—making snap judgments, checking up on their consistency, modifying them, again checking on consistency etc—leads ultimately to a bona fide, prior distribution.

Ironically, this story reverses the philosophical position of naive Bayesians. Savage's rational agents do not begin with a prior from which they deduce their posteriors. They begin by guessing at posteriors and then massage their guesses until consistency has been achieved. The prior is then derived from the system of massaged posteriors. Rational agents are therefore not somehow *endowed* with a prior. They have to work hard at *constructing* a prior (Binmore [8, pp. 128–134]).

Constructing a prior in this way would be impossible in large complicated worlds because it requires being able to look ahead and figure out what one would believe under all possible future contingencies. As Savage [67, p. 16] puts it, you must always be able to look before you leap and never have to delay considering bridges until they have to be crossed.

It is easy to understand why Bayesians who work in macroeconomics or finance find this exposition of Savage's views unwelcome, but surely a game as classically defined counts as a small world in Savage's sense? My own view is that the answer depends on whether or not the reasoning processes of the players are excluded from the formal model. If they are, then I see no conceptual difficulty in applying Bayesian decision theory to a classical game to come up with a (possibly refined) Nash equilibrium, provided that one does not need to go beyond the classical notion of a mixed strategy.

3 Ambiguity or Uncertainty?

The distinction between subjective and epistemic probabilities mentioned in the previous section is not recognized by those naive Bayesians who treat Savage's

theory as though it were a solution to the problem of scientific induction.⁵ All one ever needs to do according to this position is to name a prior and then deal with all new information by updating this prior using Bayes' rule.

Where are such rational priors to be found?⁶ Laplace [55] advocated the principle of insufficient reason (also called the indifference principle). The Harsanyi doctrine is that different rational players independently placed in a situation of complete ignorance will necessarily formulate the same common prior (Harsanyi [44]). Others take this position further by arguing that the complete ignorance assumption implies, for example, that the rational prior will maximize entropy (Jaynes and Bretthorst [47]). We are then as far from Savage's view on constructing priors as it is possible to be. Instead of using all potentially available information in the small world to be studied in formulating a prior, one treats all such information as irrelevant.

The ambiguity approach to the problem of Knightian uncertainty replaces the single prior of the preceding discussion by a (usually convex) set of priors. Alice doesn't know enough to come up with a single prior and so she rests content with eliminating those priors that she does know enough to exclude.

Such a formulation in terms of multiple priors seems an inevitable response to the problem of Knightian uncertainty when probabilities are taken to be epistemic, but there seems no particular reason always to model the problem in this way when the probabilities are subjective. Moreover, if one follows Savage's approach to constructing priors but accepts that Alice may be too ignorant to bring the massaging process to a determinate conclusion, then there is no particular reason why one should be working with priors at all. To emphasize the importance of not excluding this possibility, we shall speak of *uncertainty* rather than *ambiguity* when we do not want to take for granted that it is only Alice's ignorance that prevents her being able to assign a probability to everything.

The need to widen our horizons from ambiguity to uncertainty is defended in two ways. Section 6 expresses doubts about using the ambiguity approach when seeking to enlarge the class of mixed strategies. Section 8 casts doubt on the

⁵Philosophers distinguish objective probabilities (or chances) from epistemic probabilities (or credences). They link the two concepts using Lewis's [58] "principal principle", which says that the credence for an event whose chance is known should equal that chance. No consensus seems to exist on where subjective probabilities fit into this scheme. Do they count as chances or credences, or something in between? Only the third option seems viable to me.

⁶Bayesian statisticians use priors that have been found to work well in similar applications, but they usually do not claim that such priors are the one-and-only rational prior. Instead, they very reasonably defend their methodology against classical statistics on empirical grounds.

adequacy of the ambiguity approach in general.

4 Ellsberg Games?

The Ellsberg paradox lies at the heart of ambiguity research (Ellsberg [27, 26]). We use the version in which an urn contains 10 red balls and 20 balls that are black or white in an unknown proportion. The ambiguity approach fits this scenario without difficulty because each ball then has a specific objective probability of being drawn, and the problem is simply that we do not know what these objective probabilities are in the case of the black and white balls. All we can say for sure about these ambiguous probabilities is that they lie between a lower probability of 0 and an upper probability of $\frac{2}{3}$.

Ellsberg suggested that decision-makers asked to express preferences between various reward schedules that depend on which ball is drawn at random from the urn would not honor the postulates of Bayesian decision theory. Many experiments have confirmed this conjecture in a variety of settings. As a result, it is now regarded as a stylized fact that people usually have an inbuilt aversion to ambiguity. (See the annotated bibliography of Wakker [78]). One can argue that this result is irrelevant to rational choice theory for the same reason that experiments showing that people often assert that $7 \times 8 = 54$ is irrelevant to the fact that $7 \times 8 = 56$, but the consensus seems to be in favor of modifying Savage's theory so that ambiguity neutrality ceases to be a consequence of rationality.

In applying such work in game theory, the focus of attention has been on the approach of Gilboa and Schmeidler [33] mentioned in Section 1. Mukerji and Tallon [63] review the issues that arise and so it is enough to observe that one can contemplate introducing uncertainty aversion into a game at three levels:

1. The level of chance moves;
2. The level of beliefs about other players' strategy choices;
3. The level of strategies.

Much of the literature is devoted to the second level,⁷ which is inspired by the observation that mixed equilibria need to be interpreted as equilibria in beliefs about the strategies of the players' opponents (because players have no more reason to play a mixed strategy in equilibrium than they have to play any of the pure strategies to which it assigns positive probability). However, the focus of attention in this paper is the third level in which the strategies themselves are

⁷For example, Dow and Werlang [22], Eichberger *et al* [24, 25], Lo [59]. Marinacci [61].

modeled as being ambiguous.

5 Ellsberg Strategies?

The literature stemming from the Ellsberg paradox does not usually distinguish between large worlds and small worlds. In the kind of large-world settings in which most everyday decisions are made, some degree of uncertainty aversion certainly seems very intuitive. However, the laboratory evidence is not so solid as is commonly assumed. For example, when the same Ellsberg problem is expressed in losses rather than gains, experimenters often report ambiguity-seeking behavior rather than ambiguity-averse behavior (Wakker [77, p. 354]).⁸

What of Ellsberg experiments framed with a view to applications in small-world settings so that the possibility of ambiguous interpretations of the experimental context are minimized. Several recent experiments report very little ambiguity aversion in such circumstances (Binmore *et al* [10], Charness *et al* [16], Stahl [71], and Voorhoeve *et al* [75].) Savage would not have been surprised at this conclusion because the Ellsberg paradox is set in a world that is about as small as a world can be.

Where does this leave the “Ellsberg strategies” championed by Riedel and Sass [66]? They argue that players in a game could enlarge the set of mixed strategies by using the Ellsberg set-up. I see three problems:

1. I am one of many who think it reasonable to seek to supplement Bayesian decision theory with a theory of rational choice better adapted to large-world applications (Binmore [9]), but the proposed application of the Ellsberg paradox to implementing ambiguous strategies in rational game theory would be a small-world application to which Bayesian decision theory ought to apply.

2. The evidence that people exhibit a substantial level of ambiguity aversion in Ellsberg experiments is not robust. The evidence that subjects will decide among multiple priors using the maximin criterion is hardly there at all. It is not even clear why behavioral data is relevant to a rational theory of games.

3. The fact that Ellsberg strategies are rooted in an ambiguity approach to

⁸The Ellsberg paradox isn't the only situation in which a stylized behavioral fact is accepted uncritically. The endowment effect is sometimes said to be the most robust result in experimental economics, but half of all papers find no effect at all (Plott and Zeiler [65]). Even the Allais paradox isn't entirely solid (Blavatsky *et al* [11]). As Kahneman and Tversky [48] emphasize, how an experiment is framed can make a big difference to the observed behavior.

uncertainty is too heavy a limitation.

The next section is therefore devoted to a brief review of how one could implement what one might call “Knightian strategies”.

6 Knightian Strategies?

Von Neumann apparently said that anyone who attempts to generate random numbers by deterministic means is living in a state of sin. However, the subject of algorithmic randomness pioneered by Kolmogorov [53], Chaitin [15] and others is now well established. This computational school requires that an algorithmically generated random sequence fail all effective pattern-detecting tests (paraphrasing Calude [14]). So randomness is identified with unpredictability.

This section reviews how randomizing devices based on such theories can be used in a manner that cannot be duplicated by a mixed strategy. Binmore [8] refers to such a randomizing device as a muddling box. If such a muddling box is used in Matching Pennies and so chooses only between heads and tails, one might idealize it as an infinite sequence of 0s and 1s that represents its putative past history in selecting each alternative.

Such randomizing devices were originally introduced by Richard von Mises [74] (Ludwig’s younger brother). He says that such a device chooses 0 with probability p if the average number of 0s in the first n trials converges to p .⁹ Such a sequence is said to be random if all recursively defined subsequences also have probability p , so that no algorithmic player can profit by picking and choosing when to bet (Church [19]).

But what if the muddling box is designed not to have a probability? That is to say, the average number of 0s does not converge to a probability p but oscillates between an upper probability \bar{p} and a lower probability \underline{p} ? The ambiguity approach does not easily accommodate such a device. One might propose a (non-trivial) prior for its unknown probability and then employ Bayes’ rule to update the prior using the sequence of 0s and 1s that labels the device. But the posteriors will not eventually cluster about some particular value. They will oscillate just like the original sequence of 0s and 1s.

One might argue that such muddling boxes are only a theoretical possibility, but the device proposed by Stecher *et al* [72] for generating “ambiguity” in the

⁹It is not important here, but I favor improving the theory by strengthening the limiting requirement. This and other difficulties are discussed in Chapter 6 of Binmore [8]).

laboratory would serve as an approximation adequate for most practical purposes. One might also argue that rational players would never have a use for muddling boxes but Binmore [8, p.182] offers at least one possible application to the following problem.

What should rational players do in the Battle of the Sexes when they have no way to break the symmetry of the game? Only symmetric Nash equilibria are then viable. The only symmetric classical Nash equilibrium of the Battle of the Sexes is mixed. Since both players receive only their security payoff at this equilibrium, why they don't play their security strategies instead to make sure of this expected payoff? But when muddled strategies are allowed, the Battle of the Sexes can have a continuum of symmetric Nash equilibria, some of which pay both players strictly more than their security payoff.

The result mentioned above on muddled equilibria in the Battle of the Sexes and the clearer results described in Sections 8.2 and 9.5 differ from Bade's [6] finding that her "ambiguous act equilibria" are behaviorally equivalent to mixed-strategy equilibria in two-player games, and "consequently that there is little hope for new insights in applied contexts". This paper comes to a different conclusion because it does not shackle the players to the maximin criterion in evaluating the use of muddled strategies. The rest of the paper examines alternatives to this standard approach.

7 Upper and Lower Probabilities

A decision problem is a function $D : A \times B \rightarrow C$ in which A is a set of feasible actions, B is a set of states of the world, and C is a set of possible consequences. We always assume that the state space is a probability space (B, \mathcal{M}, p) , in which it does not matter whether \mathcal{M} is taken to be an algebra or a sigma algebra because the state space B is assumed to be finite. Similarly the probability measure p defined on \mathcal{M} may be finitely or countably additive.

We say that an event M in \mathcal{M} is *measured* because it comes with a measure $p(M)$. A subset E of B that is not in \mathcal{M} may or may not be measurable, in that the measure p may possibly admit an extension to an algebra \mathcal{N} larger than \mathcal{M} . If multiple extensions are possible, then the value of $p(E)$ will be ambiguous. However, we assume that no such extension has yet been identified, and so E remains unmeasured.

Inner and outer measure. The outer measure $\bar{p}(E)$ of a set E is the infimum of the measures $p(M)$ of all measured supersets M of E . Its inner measure $\underline{p}(E)$ is the supremum of the measures $p(M)$ of all measured subsets M of E . One can follow Lebesgue by extending a measure p to sets E on which $\underline{p}(E) = \bar{p}(E)$ in the obvious way. We assume that this extension has already been done, so that a set is measured if and only if its inner and outer measure are equal.

It follows immediately from their definitions that \bar{p} is subadditive and \underline{p} is superadditive. Also

$$\underline{p}(E) + \bar{p}(\sim E) = 1. \quad (3)$$

Following Binmore [8], Good [39], Halpern and Fagin [41], Suppes [73], and others, we identify $\bar{p}(E)$ with the upper probability of E and $\underline{p}(E)$ with its lower probability. Such an identification is only new to the extent that inner and outer measures are mentioned explicitly. For example, in the Ellsberg paradox, the algebra of measured sets is $\{\emptyset, R, \sim R, B\}$, where R is the event that a red ball is drawn. The largest measured subset of the event W that a white ball is drawn is \emptyset , and so $\underline{p}(W) = p(\emptyset) = 0$. The smallest measured superset of W is $\sim R$, and so $\bar{p}(W) = p(\sim R) = \frac{2}{3}$.

Independence. The literature contains at least four proposals on how to define independence when events are unmeasured: Bade [5, 6], Brandenburger *et al* [12], Klibanoff [50]. This paper evades the issue by restricting attention to cases where the question of how to define independence is hopefully uncontroversial.

Two measured events M and N are independent if and only if $p(M \cap N) = p(M)p(N)$. We shall only attempt to extend this notion to sets of the form $E \times F$. This notation is intended to imply that E is a subset of a state space Y and F of a state space Z . The state space B is then understood to admit the factorization $B = Y \times Z$, and our probability measure p is understood to be the product measure derived from probability measures defined on Y and Z . To write $p(E \times F) = p(E)p(F)$ when E and F are measured is therefore the same as writing $p(\{E \times Z\} \cap \{Y \times F\}) = p(E)p(F)$.

Attention is restricted to events of the form $E \times F$ to ensure that

$$\underline{p}(E \times F) = \underline{p}(E)\underline{p}(F) \quad (4)$$

$$\bar{p}(E \times F) = \bar{p}(E)\bar{p}(F). \quad (5)$$

Gambles. The set of gambles to be studied will be denoted by G . A gamble \mathbf{G} in this set is defined by

$$\mathbf{G} = \begin{array}{|c|c|c|c|c|} \hline \mathcal{P}_1 & \mathcal{P}_2 & \mathcal{P}_3 & \cdots & \mathcal{P}_m \\ \hline E_1 & E_2 & E_3 & \cdots & E_m \\ \hline \end{array}, \quad (6)$$

where $\mathcal{E} = \{E_1, E_2, \dots, E_m\}$ is a partition of the belief space B , and the prize \mathcal{P}_i is understood to result when the event E_i occurs. We assume the existence of best and worst prizes, \mathcal{W} and \mathcal{L} . A simple gamble \mathbf{S} can then be defined by

$$\mathbf{S} = \begin{array}{|c|c|} \hline \mathcal{L} & \mathcal{W} \\ \hline \sim E & E \\ \hline \end{array}, \quad (7)$$

where $\sim E$ is the complement of the set E in the state space B .

Lotteries. A lottery is a gamble in which all the events E_i that determine which prize will be received are measured. In this paper, the probabilities in the lotteries used can always be taken to be objective. We always assume that each new lottery is independent of everything else in the model.

8 Surrogate Probability

A popular approach to dealing with ambiguity problems is to propose an extension π of the probability measure p to unmeasured sets. We refer to such an extension as a surrogate probability. Schmeidler's [68] non-additive probabilities are the leading example.

Our aim in this section is to explore the implications of extending the probability measure p defined on an algebra \mathcal{M} of measured subsets of the state space B to a larger algebra \mathcal{N} using the idea of a surrogate probability π . We begin by following Anscombe and Aumann [1] in making the following postulate.

Postulate 1: *Some version of the Von Neumann and Morgenstern theory of rational decision under risk applies to lotteries with prizes in the set G of gambles constructed from events in \mathcal{N} .*

Gilboa and Schmeidler [33] are widely followed in regarding ambiguity aversion as the requirement that one can have $(1-p)f + pg \prec f$ even though $f \sim g$, where

f and g are gambles and $(1 - p)f + pg$ is the lottery in which f occurs with probability $1 - p$ and g with probability p . Postulate 1 excludes this possibility, because it implies that $(1 - p)f + pg \sim f$ when $f \sim g$. But we shall find that one can sometimes retain the possibility of ambiguity aversion without the need to deny the rationality of the Von Neumann and Morgenstern postulates when applied to objective lotteries with gambles as prizes.

To assume Postulate 1 is not to deny that there are sometimes good *behavioral* reasons for adopting Gilboa and Schmeidler's approach to ambiguity aversion. Which of the various available theories one adopts would seem to depend on the context in which it is to be applied. The theory outlined in this paper is a strictly *rational* theory. Halevy's [40] work suggests that only subjects who understand how to compound objective lotteries are likely to be rational enough for such a rational theory to apply to them in real life.

Definition. *Postulate 1 ensures the existence of a Von Neumann and Morgenstern utility function $u : G \rightarrow \mathbb{R}$. In particular, the simple gamble S of (7) has a Von Neumann and Morgenstern utility $u(S)$. We can therefore define a surrogate probability (sometimes called a matching probability) by writing*

$$\pi(E) = u(S).$$

To avoid the kind of irrationalities made fun of in Aesop's fables, one might insist at this stage that $\pi(E)$ should depend only on the properties of B as a probability space, but a lot more than this will be assumed in Postulate 4.

8.1 Ambiguity aversion?

Discussions of imprecise probabilities are often framed as though the reason for the imprecision necessarily lies in there being some ambiguity over the precise probability to be attached to an event. If this framing were always appropriate, there would be a case for insisting on the next postulate.¹⁰

Postulate 2: *If a gamble has the same Von Neumann and Morgenstern utility whatever probabilities are assigned to unmeasured events, then it should be assigned this utility when probabilities have not been assigned to these events.*

¹⁰The ambiguity interpretation fits examples like the Ellsberg paradox, but it cannot be maintained in examples like the Banach-Tarski paradox (Binmore [9]).

Theorem 1: *Either all events are measured, or else Postulates 1 and 2 imply that a surrogate probability π can satisfy the arithmetic Hurwicz criterion only when $\alpha = \frac{1}{2}$.*

Proof.¹¹ Consider a lottery in which two gambles are each available with probability $\frac{1}{2}$. In the first gamble, the decision-maker wins if the event E occurs and loses otherwise. In the second gamble, the decision-maker loses if E occurs and wins otherwise. If the utility of losing is zero and the utility of winning is 1, then the expected utility of the lottery is $\frac{1}{2}$ no matter what probability is attributed to the event E . But when π is given by the arithmetic Hurwicz criterion, we can use Postulate 1 and (3) to calculate the utility of the lottery as

$$\frac{1}{2}\{(1 - \alpha)p + \alpha P\} + \frac{1}{2}\{(1 - \alpha)(1 - P) + \alpha(1 - p)\},$$

where $p = \underline{p}(E)$ and $P = \overline{p}(E)$. It follows from Postulate 2 that

$$\frac{1}{2} = \frac{1}{2} + (1 - 2\alpha)(p - P),$$

and therefore either $p = P$ or $\alpha = \frac{1}{2}$.

8.2 Game Theory Application

Section 6 mentions the existence of Nash equilibria in muddled strategies that Pareto-dominate the mixed equilibrium of the Battle of the Sexes, but the presence in this case of two Pareto-efficient Nash equilibria in pure strategies is a distraction. We therefore obtain a similar result for a 2×2 game with a unique mixed Nash equilibrium in which the players' equilibrium outcomes and their security payoffs are the same. Such a game is necessarily asymmetric.

Without any correlate of Savage's Sure-Thing Principle (like the Independence Axiom), there is a problem about interpreting the payoffs in a game (Binmore [9]). We evade this difficulty here by interpreting the payoffs to be probabilities that players will win in simple lotteries with the best and worst prizes fixed. However, the payoffs in the strategic form of Figure 1 have been multiplied by a suitable positive factor to eliminate unnecessary fractional quantities.

In the unique classical equilibrium of the game shown in Figure 1, Row plays *down* with probability $p = \frac{3}{5}$ and Column plays *right* with probability $q = \frac{2}{5}$. At this equilibrium, Row gets only his security payoff of $14/5 = 2.8$ and Column

¹¹Theorem 3 of Binmore [9] derives a stronger conclusion from a stronger hypothesis.

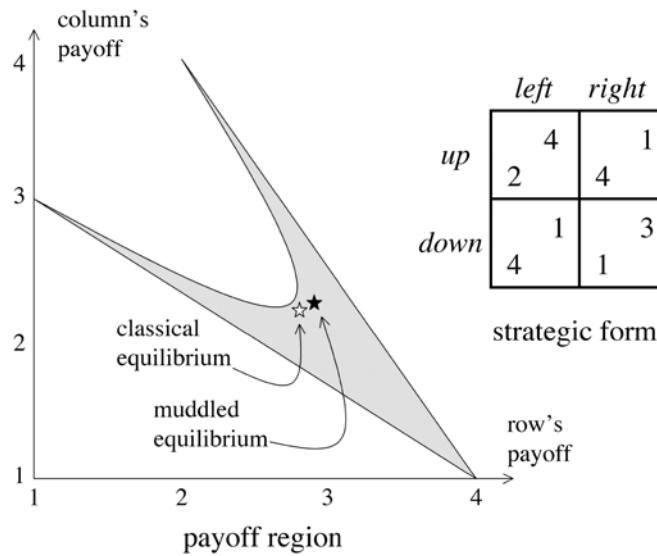


Figure 1: A game with interesting muddled equilibria. The row player's payoffs are in the SW of each cell and the column player's in the NE. The payoffs should be divided by $D \geq 4$ to allow their being interpreted as the probability of winning in a simple lottery. The payoff region shows all outcomes that can be achieved if the players use independent mixed strategies. For both the arithmetic and geometric Hurwicz criteria with $\alpha = \frac{1}{2}$, there is a 2-D continuum of muddled equilibria that Pareto-dominate the unique classical equilibrium.

only her security payoff of $11/5 = 2.2$. To secure these expected payoffs, Row must play *down* with probability $\frac{2}{5}$ and Column must play *right* with probability $\frac{3}{5}$. There is no problem with such a conclusion in an evolutionary context, but in a purely rational context one may reasonably ask what motive the players have to play their equilibrium strategies rather than their security strategies.

Appendix 1 outlines a solution to this conundrum in the case when the outcomes that result from the use of muddled strategies are evaluated using the arithmetic Hurwicz criterion with $\alpha = \frac{1}{2}$. Many muddled Nash equilibria then exist, all of which (necessarily) pay the players at least their security level (which remains unchanged when muddled strategies are allowed).

In particular, there is a muddled equilibrium in which the event that Row plays *down* has lower probability $\underline{p} = \frac{1}{10}$ and upper probability $\bar{p} = 1$, and the event that Column plays *right* has lower probability $\underline{q} = 0$ and upper probability $\bar{q} = \frac{3}{5}$. The equilibrium outcome is approximately $(3.10, 2.35)$, and so improves Row's security payoff by about 11% and Column's payoff by about 7%.

Bade’s [6] work makes it unsurprising that similar results cannot be obtained using the arithmetic Hurwicz criterion with $\alpha = 0$. Appendix 2 explains why not. However, it does not follow that the presence of ambiguity aversion is incompatible with the existence of Pareto-improving equilibria, as we show in Section 9.5 and Appendix 3.

9 Geometric Hurwicz Criterion

If ambiguity aversion is to be captured by using the arithmetic Hurwicz criterion π_A with $\alpha < \frac{1}{2}$, Theorem 1 shows that it is necessary to give up Postulate 2. However, some level of ambiguity aversion is built into the geometric Hurwicz criterion in virtue of the inequality of the arithmetic and geometric means, which implies that

$$\pi_G(E) < \pi_A(E),$$

provided that $\underline{p}(E) < \bar{p}(E)$. This variety of ambiguity aversion survives even when $\alpha = \frac{1}{2}$. But the geometric Hurwicz criterion cannot satisfy Postulate 2 at all. Its inbuilt ambiguity aversion must therefore be understood as an intrinsic aversion to imprecision rather than an instrumental response to the potential ill-effects of a lack of knowledge.

One can defend the arithmetic Hurwicz criterion as a minimal extension of Bayesian decision theory (Binmore [9]). The rest of this section presents a similar defense of the geometric Hurwicz criterion with a different interpretation of minimal.

9.1 Multiplicative Property

The essence of the argument is that since the values of a surrogate probability π substitute for probabilities—the probabilities of winning in certain simple gambles—then π should satisfy the multiplicative rule for computing the probability that two independent events E and F will both occur. That is to say, although the additive property of probabilities cannot be extended to π , we may still hope to extend the multiplicative rule as proposed in the following postulate, the plausibility of which is considered in Section 9.2.

$$\text{Postulate 3: } \quad \pi(E \times F) = \pi(E) \times \pi(F). \quad (8)$$

It follows immediately from (4) and (5) that the geometric Hurwicz criterion satisfies Postulate 3. For a converse result, we need to insist that the surrogate probability $\pi(E)$ depends only on the upper and lower probabilities of the event E , an assumption that seems unavoidable unless information other than that built into the probability measure p is available. This assumption is incorporated in the next postulate, in which $D = \{(x, y) : 0 < x \leq y \leq 1\}$.¹²

Postulate 4: *There exists a continuously differentiable function $v : D \rightarrow \mathbb{R}$ with $v(p, p) = p$ such that for all events E in \mathcal{N} ,*

$$\pi(E) = v(\underline{p}(E), \bar{p}(E)). \quad (9)$$

Theorem 2. *Postulates 3 and 4 imply the geometric Hurwicz criterion:*

$$v(\underline{p}, \bar{p}) = \underline{p}^{1-\alpha} \bar{p}^\alpha.$$

for some α with $0 \leq \alpha \leq 1$.

Proof In view of (4) and (5), Postulates 3 and 4 imply that

$$v(pq, PQ) = v(p, P) v(q, Q), \quad (10)$$

where P and p are the upper and lower probabilities of E , and Q and q are the upper and lower probabilities of F . By Postulate 4, it follows that

$$qv_1(pq, PQ) = v_1(p, P)v(q, Q).$$

Take $p = P = 1$ in this equation:

$$qv_1(q, Q) = v_1(1, 1)v(q, Q).$$

Integrate this differential equation with respect to q :

$$v(q, Q) = B(Q)q^\beta,$$

where $\beta = v_1(1, 1)$. Similarly,

$$v(q, Q) = A(q)Q^\alpha,$$

where $\alpha = v_2(1, 1)$. Comparing these two equations for $v(q, Q)$, we find that

$$v(q, Q) = Cq^\beta Q^\alpha,$$

for some absolute constant C . But $v(q, q) = q$, and so $C = 1$ and $\alpha + \beta = 1$.

¹²Postulate 4 with D replaced by $D = \{(x, y) : 0 \leq x \leq y \leq 1\}$ leads to the arithmetic Hurwicz criterion (Binmore [9]).

9.2 Subjectively Independent Events

When is Postulate 3 a reasonable hypothesis? To address this question, we ask what assumptions about a decision-maker's preferences are needed to justify, not only Postulate 3, but the standard formula for multiplying the probabilities of independent measured events when the probabilities are subjective. For this discussion, we assume Postulate 4, although a much weaker assumption would be adequate.

Separating preferences. We begin by reviewing a result of Keeney and Raiffa [49] that applies when the consequence space C can be factored so that $C = C_1 \times C_2$. The prizes in C then take the form $(\mathcal{P}, \mathcal{Q})$, where \mathcal{P} is a prize in C_1 and \mathcal{Q} is a prize in C_2 . A preference relation \preceq on C evaluates C_1 and C_2 *separately* if and only if

$$\begin{aligned} (\mathcal{P}, \mathcal{Q}) \prec (\mathcal{P}, \mathcal{Q}') &\text{ implies } (\mathcal{P}', \mathcal{Q}) \preceq (\mathcal{P}', \mathcal{Q}') ; \\ (\mathcal{P}, \mathcal{Q}) \prec (\mathcal{P}', \mathcal{Q}) &\text{ implies } (\mathcal{P}, \mathcal{Q}') \preceq (\mathcal{P}', \mathcal{Q}') . \end{aligned}$$

If the consequence spaces C , C_1 and C_2 are respectively replaced in this definition by the sets of all independent lotteries over these outcome spaces, the separation requirement is surprisingly strong (Binmore [8, p. 47]).

If \preceq can be represented by a Von Neumann and Morgenstern utility function $u : C \rightarrow \mathbb{R}$, then

$$u = A u_1 u_2 + B u_1 (1 - u_2) + C u_2 (1 - u_1) + D (1 - u_1) (1 - u_2) , \quad (11)$$

where the functions $u_1 : C_1 \rightarrow \mathbb{R}$ and $u_2 : C_2 \rightarrow \mathbb{R}$ can be regarded as normalized Von Neumann and Morgenstern utility functions on C_1 and C_2 . The constants in (11) are $A = u(\mathcal{W}_1, \mathcal{W}_2)$, $B = u(\mathcal{W}_1, \mathcal{L}_2)$, $C = u(\mathcal{L}_1, \mathcal{W}_2)$ and $D = u(\mathcal{L}_1, \mathcal{L}_2)$.

Defining subjective independence. Instead of taking Kolmogorov's definition of independence for granted, we begin by assuming that the state space B can be factored so that $B = B_1 \times B_2$, and ask what it might mean to say that the decision-maker regards all events E_1 in B_1 as independent of all events E_2 in B_2 .

To test for independence, we need to introduce a suitable consequence space C that can be similarly factored so that $C = C_1 \times C_2$. It is enough if $C_i = \{\mathcal{L}_i, \mathcal{W}_i\}$. However, it is essential that

$$u(\mathcal{W}_1, \mathcal{L}_2) = u(\mathcal{L}_1, \mathcal{W}_2) = u(\mathcal{L}_1, \mathcal{L}_2) = 0 . \quad (12)$$

so that losing in either C_1 or C_2 is equivalent to losing in both.

Let J_i be the space of simple gambles \mathbf{H}_i in which \mathcal{W}_i is obtained if and only if an event E_i in B_i occurs. We can then construct the space $J = J_1 \times J_2$ of pairs of simple gambles of the form $\mathbf{H} = (\mathbf{H}_1, \mathbf{H}_2)$. In view of (12), we identify \mathbf{H} with the simple gamble in which the decision-maker wins $\mathcal{W} = (\mathcal{W}_1, \mathcal{W}_2)$ if and only the event $E_1 \times E_2$ occurs (and gets $\mathcal{L} = (\mathcal{L}_1, \mathcal{L}_2)$ otherwise).

Definition. All events E_1 in B_1 are subjectively independent of all events E_2 in B_2 if and only if the preference relation \preceq evaluates J_1 and J_2 separately.

Apply equation (11) to the case when C , C_1 and C_2 are respectively replaced by J , J_1 and J_2 . It then reduces to $u = u_1 u_2$, which we can rewrite as

$$\pi(E_1 \times E_2) = \pi(E_1) \times \pi(E_2),$$

in which the use of the same function π when dealing with the spaces B_1 and B_2 as with B is justified by Postulate 4.

9.3 Bayesian Updating

The defining formula $p(E \cap F) = p(E | F) p(F)$ for a conditional probability is multiplicative, and hence treats the event F and the “conditional event” $E|F$ as independent. Extending the preceding reasoning to this new situation, we have

$$\pi(E \cap F) = \pi(E | F) \pi(F),$$

even for unmeasured events, and so Bayes’ rule continues to apply when updating the surrogate probability π . In view of Theorem 2, this extended version of Bayes’ rule would seem to make sense only in the case of the geometric Hurwicz criterion, in which case

$$\pi(E | F) = \frac{\pi(E \cap F)}{\pi(F)} = \left\{ \frac{\underline{p}(E \cap F)}{\underline{p}(F)} \right\}^{1-\alpha} \left\{ \frac{\overline{p}(E \cap F)}{\overline{p}(F)} \right\}^{\alpha},$$

provided that $\pi(F) \neq 0$. One therefore updates upper and lower probabilities when using the geometric Hurwicz criterion exactly as if they were ordinary probabilities.

9.4 Arithmetic Approximation

Recall that a surrogate probability is defined in Section 8 as the utility of a certain simple lottery (7). It therefore depends on the choice of \mathcal{L} and \mathcal{W} (which must be kept fixed in any particular analysis). When working with a multiplicative rule like the geometric Hurwicz criterion, the choice of \mathcal{L} plays a special role because it corresponds to events E with $\pi(E) = 0$. As a result, the classical assumption that it is inconsequential to add a fixed amount b to a player's payoffs in a game ceases to be true. (The same is sometimes true in evolutionary applications of game theory.)

In practice, one would normally identify \mathcal{L} with some hellish event remote from our experience, and then replace a payoff x in a game like that of Figure 1 by $b + \epsilon x$, where $\epsilon > 0$ is small compared with b . It would then be difficult to distinguish the geometric Hurwicz criterion from the arithmetic Hurwicz criterion in an experiment because

$$\{b + \epsilon \underline{p}\}^{1-\alpha} \{b + \epsilon \bar{p}\}^\alpha = (1 - \alpha)\{b + \epsilon \underline{p}\} + \alpha\{b + \epsilon \bar{p}\} + O(\epsilon^2).$$

9.5 Game Theory Application Again

Section 8.2 reports the existence of muddled equilibria for the game of Figure 1 that Pareto-dominate the unique classical equilibrium when muddled outcomes are evaluated using the arithmetic Hurwicz criterion with $\alpha = \frac{1}{2}$ (but not with $\alpha = 0$). This conclusion is reached without attributing any ambiguity aversion to the players, with the result that the best replies to the Pareto-improving equilibrium strategies are necessarily pure. However, neither restriction is necessary. Appendix 3 shows that Pareto-improving muddled equilibria also exist when the arithmetic Hurwicz criterion is replaced by the corresponding geometric criterion.

In particular, there is a muddled equilibrium in which the event that Row plays *down* has lower probability $\underline{p} = 0.09$ and upper probability $\bar{p} = 0.85$, and the event that Column plays *right* has lower probability $\underline{q} = 0.15$ and upper probability $\bar{q} = 0.58$. The equilibrium outcome is $(3.05, 2.33)$, and so improves Row's security payoff by 9% and Column's security payoff by 6%.

10 Conclusion

This paper explores the possibility of introducing Knightian uncertainty into the rational analysis of games. It argues that the ambiguity approach to uncertainty

is inadequate for all purposes. Where it is adequate, there are also grounds for rejecting the maximin criterion for evaluating ambiguous events in favor of the ambiguity-neutral Hurwicz criterion. However, uncertainty aversion can be retained by replacing the orthodox Hurwicz criterion by a version in which the weighted arithmetic mean of upper and lower probabilities is replaced by the corresponding geometric mean. Foundational arguments are offered in support of such an approach.

The paper continues by arguing that the Ellsberg experimental set-up is inadequate for implementing Knightian strategies, and proposes instead appealing to the theory of algorithmic randomization to realize “muddled strategies” that expand the set of mixed strategies in much the same way that mixed strategies expand the set of pure strategies. An example is given of a simple game with a unique classical equilibrium that admits Pareto-improving muddled equilibria.

References

- [1] F. Anscombe and R. Aumann. A definition of subjective probability. *Annals of Mathematical Statistics*, 34:199–205, 1963.
- [2] K. Arrow and L. Hurwicz. An optimality criterion for decision-making under ignorance. In C. Carter and J. Ford, editors, *Uncertainty and Expectations in Economics*. Basil Blackwell, Oxford, 1972.
- [3] R. Aumann and M. Maschler. Some thoughts on the minimax principle. *Management Science*, 18:54–63, 1972.
- [4] Y. Azrieli and R. Teper. Uncertainty aversion and equilibrium existence in games with incomplete information. *Games and Economic Behavior*, 73:310–317, 2011.
- [5] S. Bade. Stochastic independence with maxmin expected utilities. (Pennsylvania State University working paper), 2008.
- [6] S. Bade. Ambiguous act equilibria. *Games and Economic Behavior*, 71:246–260, 2011.
- [7] J. Berger. *Statistical Decision Theory and Bayesian Analysis*. Springer, Berlin, 1985.
- [8] K. Binmore. *Rational Decisions*. Princeton University Press, Princeton, 2009.
- [9] K. Binmore. A minimal extension of Bayesian decision theory. *Theory and Decision*, 86:25–41, 2016.

- [10] K. Binmore, L. Stewart, and A. Voorhoeve. How much ambiguity aversion? Finding indifference between Ellsberg's risky and ambiguous bets. *Journal of Risk and Uncertainty*, 45:215–238, 2012.
- [11] P. Blavatsky, A. Ortmann, and V. Panchenko. Now you see it. now you don't: How to make the Allais paradox appear or disappear. (School of Economics, UNSW Business School), 2015.
- [12] L. Blume, A. Brandenburger, and E. Dekel. Lexicographic probabilities and choice under uncertainty. *Econometrica*, 59:61–80, 1991.
- [13] J. Breese and K. Fertig. Decision making with interval influence diagrams. In P. Bonissone, M. Henrion, L. Kanal, and J. Lenner, editors, *Uncertainty in Artificial Intelligence*. Elsevier, Dordrecht, 1991.
- [14] C. Calude. *Information and Randomness*. Springer, Berlin, 1998.
- [15] G. Chaitin. *Exploring Randomness*. Springer, Berlin, 2001.
- [16] D. Charness, E. Karni, and D. Levin. Ambiguity attitudes and social interactions: An experimental investigation. *Journal of Risk and Uncertainty*, 46:1–25, 2013.
- [17] H. Chernoff. Rational selection of decision functions. *Econometrica*, 22:422–443, 1954.
- [18] L. Chrisman. Incremental conditioning of upper and lower probabilities. *International Journal of Approximate Reasoning*, 13:1–25, 1995.
- [19] A. Church. On the concept of a random sequence. *Bulletin of the American Mathematical Society*, 46:130–135, 1940.
- [20] B. de Finetti. *Theory of Probability, Volume II*. Wiley, London, 1975.
- [21] A. Dempster. Upper and lower probabilities induced by a multivalued mapping. *Annals of Mathematical Statistics*, 38:325–339, 1967.
- [22] J. Dow and C. Werlang. Nash equilibrium under Knightian uncertainty: Breaking down backward induction. *Journal of Economic Theory*, 64:305–324, 1994.
- [23] J. Earman. *Bayes or Bust?* MIT Press, Cambridge, MA, 1992.
- [24] J. Eichberger and D. Kelsey. Non-additive beliefs and strategic equilibria. *Games and Economic Behavior*, 30:355–379, 2000.
- [25] J. Eichberger, D. Kelsey, and B. Schipper. Ambiguity and social interaction. *Oxford Economic Papers*, 61:355–379, 2009.

- [26] D. Ellsberg. The crude analysis of strategy choices. *American Economic Review*, 51:472–478, 1961.
- [27] D. Ellsberg. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75:643–669, 1961.
- [28] K. Fertig and J. J. Breese. Interval influence diagrams. In M. Henrion, P. Shachter, L. Kanal, and J. Lenner, editors, *Uncertainty in Artificial Intelligence*. Elsevier, Dordrecht, 1990.
- [29] T. Fine. *Theories of Probability*. Academic Press, London, 1973.
- [30] T. Fine. Lower probability models for uncertainty and nondeterministic processes. *Journal of Statistical Planning and Inference*, 20:389–441, 1988.
- [31] P. Fishburn. *Utility Theory for Decision Making*. Wiley, New York, 1970.
- [32] P. Fishburn. *Foundations of Expected Utility*. Kluwer, Amsterdam, 1982.
- [33] P. Gardenfors and N. Sahlin. Unreliable probabilities, risk taking and decision making. *Synthese*, 53:361–386, 1982.
- [34] P. Ghirardato, F. Maccheroni, and M. Marinacci. Ambiguity from the differential viewpoint. Social Science Working Paper 1130, 2002.
- [35] P. Ghirardato and M. Marinacci. Ambiguity made precise: A comparative foundation. *Journal of Economic Theory*, 102:251–289, 2002.
- [36] I. Gilboa. *Uncertainty in Economic Theory: Essays in Honor of David Schmeidler's 65th Birthday*. Routledge, London, 2004.
- [37] I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1989.
- [38] F. Giron and S. Rios. Quasi-Bayesian behavior; a more realistic approach to decision-making? In D. Lindley J. Bernardo, J. DeGroot and A. Smith, editors, *Bayesian Statistics*. University Press, Valencia, Spain, 1980.
- [39] I. J. Good. *Good Thinking: The Foundations of Probability and its Applications*. University of Minnesota Press, Minneapolis, 1983.
- [40] Y. Halevy. Ellsberg revisited: An experimental study. *Econometrica*, 75:503–506, 2007.
- [41] J. Halpern and R. Fagin. Two views of belief: Belief as generalized probability and belief as evidence. *Artificial Intelligence*, 54:275–3178, 1992.

- [42] J. Harsanyi. A general solution for finite non-cooperative games, based on risk dominance. In L. Shapley M. Dresher and A. Tucker, editors, *Advances in Game Theory*. Princeton University Press, Princeton, 1964.
- [43] J. Harsanyi. A general theory of rational behavior in game situations. *Econometrica*, 34:613–634, 1966.
- [44] J. Harsanyi. *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge University Press, Cambridge, 1977.
- [45] P. Huber. *Robust Statistics*. Wiley, New York, 1980.
- [46] L. Hurwicz. Optimality criteria for decision making under ignorance. Cowles Commission Discussion Paper, Statistics 370, 1951.
- [47] E. Jaynes and L. Bretthorst. *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge, 2003.
- [48] D. Kahneman and A. Tversky. The framing of decisions and the psychology of choice. *Science*, 211:453–458, 1981.
- [49] R. Keeney and H. Raiffa. Additive value functions. In J. Ponsard et al, editor, *Theorie de la Decision et Applications*. Fondation national pour L'Eneignement de la Gestion des Enterprises, Paris, 1975.
- [50] P. Klibanoff. Stochastically independent randomization and uncertainty aversion. *Economic Theory*, 18:605–620, 2001.
- [51] P. Klibanoff, M. Marinacci, and S. Mukerji. A smooth model of decision-making under ambiguity. *Econometrica*, 73:1849–1892, 2005.
- [52] F. Knight. *Risk, Uncertainty, and Profit*. Houghton-Mifflin, Boston, 1921.
- [53] A. Kolmogorov. Three approaches for defining the concept of 'information quantity'. *Probl. Inf. Transm. J.*, 14:3–11, 1968.
- [54] H. Kyburg. Bayesian and non-Bayesian evidential updating. *Artificial Intelligence*, 31:271–293, 1987.
- [55] P.-S. de Laplace. *A Philosophical Essay on Probabilities*. Wiley, New York, 1902. (Translated by R. Brambrough).
- [56] I. Levi. *The Enterprise of Knowledge*. MIT Press, Cambridge, MA, 1980.
- [57] I. Levi. *Hard Choices; Decision Making under Unresolved Conflict*. Cambridge University Press, Cambridge, MA, 1986.

- [58] D. Lewis. A subjectivist' guide to objective chances. In *Studies in Inductive Logic and Probability*. University of California Press, Berkeley, CA, 1980.
- [59] K. Lo. Equilibrium in beliefs under uncertainty. *Journal of Economic Theory*, 71:443–484, 1996.
- [60] R. Luce and H. Raiffa. *Games and Decisions*. Wiley, New York, 1957.
- [61] M. Marinacci. Ambiguous games. *Games and Economic Behavior*, 31:191–219, 2000.
- [62] J. Milnor. Games against Nature. In *Decision Processes*. Wiley, New York, 1954. (Edited by R. Thrall, C. Coombs, and R. Davies).
- [63] S. Mukerji and J-M. Tallon. An Overview of economic applications of David Schmeidler's models of decision-making under uncertainty. In I. Gilboa, editor, *Uncertainty in Economic Theory: A Collection of Essays in Honor of David Schmeidler's 65th birthday*. Routledge, London, 2004.
- [64] J. Pearl. On probability intervals. *International Journal of Approximate Reasoning*, 2:211–216, 1988.
- [65] C. Plott and K. Zeiler. The willingness to pay/willingness to accept gap, the “endowment effect”, subject misconceptions and experimental procedures for eliciting valuations. *American Economic Review*, 95:530–545, 2005.
- [66] F. Riedel and L. Sass. Ellsberg strategies. (working paper 452 of Institute of Mathematical Economics, Bielefeld, 2013).
- [67] L. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.
- [68] D. Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57:571–585, 1989.
- [69] T. Seidenfeld. Outline of a theory of partially ordered preferences. *Philosophical Topics*, 21:173–188, 1993.
- [70] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, 1976.
- [71] D. Stahl. Heterogeneity of ambiguity preferences. *Review of Economics and Statistics*, 96:609–617. 2014.
- [72] J. Stecher, T. Shields, and J. Dickhaut. Generating ambiguity in the laboratory. *Management Science*, 57:705–712, 2011.

- [73] P. Suppes. The measurement of belief. *Journal of the Royal Statistical Society*, 2:160–191, 1974.
- [74] R. von Mises. *Probability, Statistics, and Truth*. Allen and Unwin, London, 1928.
- [75] A. Voorhoeve, K. Binmore, A. Stefansson, and L. Stewart. Ambiguity attitudes, framing, and consistency. <http://personal.lse.ac.uk/voorhoev/>, 2015.
- [76] P. Wakker. Continuous subjective expected utility with nonadditive probabilities. *Journal of Mathematical Economics*, 18:1–27, 1989.
- [77] P. Wakker. *Prospect Theory for Risk and Ambiguity*. Cambridge University Press, Cambridge, 2010.
- [78] P. Wakker. Annotated bibliography. people.few.eur.nl/wakker/refs/webfrncs.doc, 2015.
- [79] A. Wald. *Statistical Decision Theory*. Wiley, New York, 1950.
- [80] P. Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman Hall, London, 1991.
- [81] P. Walley and T. Fine. Towards a frequentist theory of upper and lower probability. *Annals of Statistics*, 10:741–761, 1982.

Appendix 1: Ambiguity Neutral Case

This appendix outlines the proof that the game of Figure 1 admits a muddled equilibrium that is a Pareto-improvement on the unique classical equilibrium when muddled outcomes are evaluated using the arithmetic Hurwicz criterion with $\alpha = \frac{1}{2}$.

If Row plays *down* with probability p and Column plays *right* with probability q , then their respective expected payoffs are¹³

$$\pi_1(p, q) = \frac{14}{5} - 5(p - \frac{2}{5})(q - \frac{2}{5}) \quad (13)$$

$$\pi_2(p, q) = \frac{11}{5} + 5(p - \frac{3}{5})(q - \frac{3}{5}) \quad (14)$$

One can read off that there is a unique Nash equilibrium at $(p, q) = (\frac{3}{5}, \frac{2}{5})$ that pays Row $\frac{14}{5}$ and Column $\frac{11}{5}$. Also that $p = \frac{2}{5}$ secures $\frac{14}{5}$ for Row and $q = \frac{3}{5}$ secures $\frac{11}{5}$ for Column.

¹³These payoffs are to be understood as proportional to probabilities of winning in simple lotteries with fixed prizes.

To ease the algebra, we take $u = p - \frac{3}{5}$ and $v = q - \frac{2}{5}$. Writing $\pi_1(p, q) = \Pi_1(u, v)$ and $\pi_2(p, q) = \Pi_2(u, v)$, we have

$$\Pi_1(u, v) = \frac{14}{5} \left\{ 1 - \frac{25}{14} \left(u + \frac{1}{5} \right) v \right\} \quad (15)$$

$$\Pi_2(u, v) = \frac{11}{5} \left\{ 1 + \frac{25}{11} u \left(v - \frac{1}{5} \right) \right\} \quad (16)$$

A muddled strategy for Row is a pair (\underline{p}, \bar{p}) with $0 \leq \underline{p} \leq \bar{p} \leq 1$. A muddled strategy for Column is a pair (\underline{q}, \bar{q}) with $0 \leq \underline{q} \leq \bar{q} \leq 1$. We write $x = \underline{p} - \frac{3}{5}$, $X = \bar{p} - \frac{3}{5}$, $y = \underline{q} - \frac{2}{5}$, and $Y = \bar{q} - \frac{2}{5}$. The plan is to seek a muddled equilibrium with $x \leq 0 \leq X$ and $y \leq 0 \leq Y$ (New equilibria do not otherwise exist.) We then have the constraints

$$-\frac{3}{5} \leq x \leq 0 \leq X \leq \frac{2}{5} \quad (17)$$

$$-\frac{2}{5} \leq y \leq 0 \leq Y \leq \frac{3}{5} \quad (18)$$

In all the cases we consider, the best replies to a muddled strategy always include a mixed strategy. In seeking a muddled equilibrium, we can sometimes use such an optimal mixed strategy to determine the upper or lower probability of a muddled strategy that is also optimal.

We use ξ to denote a mixed strategy for Row that is a best reply to the muddled strategy (y, Y) for Column. In seeking to make (x, X) a best reply to (y, Y) , we can try $x = \xi$ or $X = \xi$. When $x = \xi$, (x, X) will continue to be a best reply to (y, Y) if $\Pi_1(X, y)$ and $\Pi_1(X, Y)$ lie between $\Pi_1(x, y)$ and $\Pi_1(x, Y)$, because Hurwicz criteria pay attention only to extreme values. Similar considerations apply when $X = \xi$, leading to the following criteria for (x, X) to be a best reply

$$\xi = x \Rightarrow \frac{x + \frac{1}{5}}{X + \frac{1}{5}} \leq \min \left\{ \frac{y}{Y}, \frac{Y}{y} \right\} \quad (\leq -1) \quad (19)$$

$$\xi = X \Rightarrow \frac{x + \frac{1}{5}}{X + \frac{1}{5}} \geq \max \left\{ \frac{y}{Y}, \frac{Y}{y} \right\} \quad (\geq -1) \quad (20)$$

Similarly, we use η to denote a mixed strategy for Column that is a best reply to the muddled strategy (x, X) for Row. The corresponding criteria to (19) and (20) are then

$$\eta = y \Rightarrow \frac{Y - \frac{1}{5}}{y - \frac{1}{5}} \geq \max \left\{ \frac{x}{X}, \frac{X}{x} \right\} \quad (\geq -1) \quad (21)$$

$$\eta = Y \Rightarrow \frac{Y - \frac{1}{5}}{y - \frac{1}{5}} \leq \min \left\{ \frac{x}{X}, \frac{X}{x} \right\} \quad (\leq -1) \quad (22)$$

The preceding analysis is the same for all Hurwicz criteria. What follows is special to the case of the arithmetic Hurwicz criterion with $\alpha = \frac{1}{2}$. In locating a mixed best reply to (y, Y) , Row then needs to find the value of $u = \xi$ subject to $-\frac{3}{5} \leq u \leq \frac{2}{5}$ that maximizes $\frac{14}{5}A_1$, where

$$A_1 = 1 - \frac{25}{28}(u + \frac{1}{5})(y + Y) \quad (23)$$

Similarly, Column needs to find the value of $v = \eta$ subject to $-\frac{2}{5} \leq v \leq \frac{3}{5}$ that maximizes $\frac{11}{5}A_2$, where

$$A_2 = 1 - \frac{25}{22}(v - \frac{1}{5})(x + X) \quad (24)$$

We recover the classical mixed equilibrium of the game by taking $x = X = 0$ and $y = Y = 0$. There are muddled analogues with $x + X = 0$ and $y + Y = 0$ in which both players get only their security payoffs. There is an isolated muddled equilibrium in which Row plays $(x, X) = (-\frac{3}{5}, \frac{1}{5})$ and Column plays $(y, Y) = (-\frac{2}{5}, \frac{2}{5})$. Column then improves on her security payoff but Row does not.

The interesting muddled equilibria arise when $\xi = X = \frac{2}{5}$ and $\eta = y = -\frac{2}{5}$. For $\xi = \frac{2}{5}$, we need $y + Y \geq 0$ in (23). For $\eta = -\frac{2}{5}$, we need $x + X \leq 0$ in (24). In addition to (17) and (18), we therefore have the additional constraints: $x \leq -\frac{2}{5}$ and $Y \leq \frac{2}{5}$. Two further constraints are obtained from (20) and (21). The requirements for an equilibrium of this type are therefore:

$$-\frac{3}{5} \leq x \leq -\frac{2}{5} \quad (25)$$

$$0 \leq Y \leq \frac{2}{5} \quad (26)$$

$$2x + 3Y \geq -\frac{2}{5} \quad (27)$$

$$3x + 2Y \leq \frac{2}{5} \quad (28)$$

The Pareto-efficient equilibria satisfy $2x + 3Y = -\frac{2}{5}$. An example is given in Section 8.2.

Appendix 2: Maximin Case

The analysis is the same as Appendix 1 until equations (23) and (24), which need to be replaced by maximin criteria. Row needs to find the value of $u = \xi$ subject to $-\frac{3}{5} \leq u \leq \frac{2}{5}$ that maximizes $\frac{14}{5}M_1$, where

$$M_1 = \begin{cases} 1 - \frac{25}{14}(u + \frac{1}{5})y, & u \leq -\frac{1}{5} \\ 1 - \frac{25}{14}(u + \frac{1}{5})Y, & u \geq -\frac{1}{5} \end{cases} \quad (29)$$

It follows that either $y = Y$ or the maximizing u is $\xi = -\frac{1}{5}$ (which is Row's security strategy). Similarly, Column needs to find the value of $v = \eta$ subject to $-\frac{2}{5} \leq v \leq \frac{3}{5}$ that maximizes $\frac{11}{5}M_2$, where

$$M_2 = \begin{cases} 1 + \frac{25}{11}x(v - \frac{1}{5}), & v \geq \frac{1}{5} \\ 1 + \frac{25}{11}X(v - \frac{1}{5}), & v \leq \frac{1}{5} \end{cases} \quad (30)$$

It follows that either $x = X$ or the maximizing v is $\eta = \frac{1}{5}$ (which is Column's security strategy).

Only the classical mixed equilibrium with $x = X = 0$ and $y = Y = 0$ is therefore viable.

Appendix 3: Geometric Case

The analysis is the same as Appendix 1 until equations (23) and (24), which need to be replaced by corresponding geometric criteria. Row needs to find the value of $u = \xi$ subject to $-\frac{3}{5} \leq u \leq \frac{2}{5}$ that maximizes $\frac{14}{5}\sqrt{G_1}$, where

$$G_1 = \{1 - \frac{25}{14}(u + \frac{1}{5})y\}\{1 - \frac{25}{14}(u + \frac{1}{5})Y\} \quad (31)$$

In an appropriate range, the value of u that maximizes G_1 is found by setting its derivative to zero. The mixed strategy for Row that maximizes G_1 therefore corresponds to $u = \xi$, where

$$\xi = -\frac{1}{5} + \frac{7}{25}\{y^{-1} + Y^{-1}\} \quad (32)$$

subject to the constraints

$$(y + \frac{7}{10})(Y + \frac{7}{10}) \leq \frac{49}{100} \quad \text{and} \quad (y - \frac{7}{15})(Y - \frac{7}{15}) \leq \frac{49}{225} \quad (33)$$

Similarly, the mixed strategy corresponding to v that maximizes Column's payoff when Row uses the muddled strategy corresponding to (x, X) is $v = \eta$, where

$$\eta = \frac{1}{5} - \frac{11}{5}\{x^{-1} + X^{-1}\} \quad (34)$$

subject to the constraints

$$(x - \frac{11}{30})(X - \frac{11}{30}) \leq \frac{121}{900} \quad \text{and} \quad (x + \frac{11}{20})(X + \frac{11}{20}) \leq \frac{121}{400} \quad (35)$$

As in Appendix 1, a muddled equilibrium might have $\xi = x$ or $x = X$, and $\eta = y$ or $\eta = Y$. However, the possibilities with $\xi = x$ or $\eta = Y$ are ruled out

because (19) implies $x + X \leq -\frac{2}{5}$ which is incompatible with (35), and (21) implies that $y + Y \geq \frac{2}{5}$ which is incompatible with (33).

We therefore seek a muddled equilibrium with $\xi = X$ and $\eta = y$. Such an equilibrium consists of a pair (x, X) for Row and a pair (y, Y) for Column such that (32) and (34) hold, subject to the constraints (17), (18), (20), (21), (33) and (35). There is a continuum of such equilibria, including the example cited in Section 9.5.