

Assessing strategic risk

R.J. Aumann and J. Dreze

Discussion Paper 2005-20

Département des Sciences Économiques
de l'Université catholique de Louvain



UCL

Assessing Strategic Risk*

R. J. Aumann[†] and J. H. Dreze[‡]

March 17, 2005

Abstract

In recent decades, the concept of subjective probability has been increasingly applied to an adversary's choices in strategic games. A careful examination reveals that the standard construction of subjective probabilities does not apply in this context. We show how the difficulty may be overcome by means of a different construction, and provide an axiomatic foundation for it.

*The authors are grateful to Sergiu Hart, Martin Meier, Jean-Francois Mertens, Bernard Walliser and Shmuel Zamir for important input. We alone remain responsible for the contents.

[†]Department of Mathematics and Center for Rationality, The Hebrew University of Jerusalem, 91904 Jerusalem, Israel; raumann@math.huji.ac.il

[‡]CORE, Université Catholique de Louvain, 34, Voie du Roman Pays, 1348 Louvain-la-Neuve, Belgium; dreze@core.ucl.ac.be

1. Introduction

Half a century ago, when decision theory and game theory were young, it was common to perceive a dichotomy between (i) *games against nature*, in which the “adversary” is a neutral “nature”—and (ii) *strategic games*, in which the adversary is an interested party or parties. Games against nature were analyzed using several criteria, most prominent being the maximization of subjective expected utility—i.e., expected utility when the probabilities assigned to nature’s moves are “subjective” or “personal,” as in Savage [1954];¹ whereas strategic games were analyzed by minimax, or, more generally, strategic equilibrium in the sense of Nash [1951]. No need was seen to reconcile or even relate the approaches, perceived as proceeding from distinct conceptual foundations.

In the ensuing years, the dichotomy gradually disappeared. It was recognized that games against nature and strategic games are in principle quite similar, and can—perhaps should—be treated similarly. Specifically, a player in a strategic game should be able to form subjective probabilities over the strategies of the other players, and then choose his own strategy so as to maximize his expected utility with respect to these subjective probabilities.

There is, however, a difficulty with applying the notion of subjective probability to strategic games. In games against nature, subjective probabilities are constructed from the the decision maker’s preferences among gambles that “stake prizes” on nature’s possible choices. In strategic games, this translates to staking prizes on strategy choices of the adversary. But that may change the incentives of the players, and in particular, the propensity to play this or that strategy. This invalidates the whole process.

In this paper, we propose a resolution of this difficulty.

The nature of the difficulty is described more precisely in the following section. In Section 3, we describe the resolution informally. Section 4 is devoted to mathematical preliminaries, Section 5 to the formal statement of our result, Section 6 to discussion, Sections 7 and 8 to proofs, and Section 9 to the literature.

2. The Difficulty

It will be useful to use the same terminology for games against nature and for strategic games. In either case, the “adversary”—be it nature or an inter-

¹Related ideas occur in Ramsey [1931], DeFinetti [1937], Dreze [1961], and Anscombe and Aumann [1963].

ested party or parties—has several alternatives, called *strategies of the adversary*. The decision maker—henceforth *protagonist*—also has several alternatives, called *strategies of the protagonist*. Together, the strategies of the adversary and of the protagonist determine the outcome of the game. Thus, each of the protagonist’s strategies may be thought of as a function from the adversary’s strategies to the possible outcomes: an “act” in the terminology of Savage [1954], a “horse lottery” in that of Anscombe and Aumann [1963] (henceforth A-A).

Before proceeding, we review and contrast the concepts of objective and subjective probability. *Objective* probabilities are associated with processes that are repeatable—and/or in which considerations of symmetry determine the probability—like coin tosses, roulette spins, dice throws, and so on. *Subjective* probabilities are associated with one-time events like elections, tomorrow’s weather, a horse race or how an adversary will play a game. Reasonable people may be expected to agree on the numerical values of objective probabilities, but may well differ on subjective probabilities. That is why subjective probabilities are often called “personal:” They are associated with a particular person i . Roughly, i ’s subjective probability for an event E is that number p such that i would as soon have a dollar with objective probability p as a dollar contingent on E . In our case, the protagonist’s subjective probability for the adversary playing some specific strategy s is that number p such that she² would as soon have a dollar with objective probability p as a dollar contingent on his playing s .

More precisely, let $\mathbf{1}$ and $\mathbf{0}$ be outcomes, with the protagonist preferring $\mathbf{1}$ to $\mathbf{0}$. For each of the adversary’s strategies s , the protagonist considers a hypothetical strategy r of her own yielding $\mathbf{1}$ if the adversary plays s , and $\mathbf{0}$ otherwise. Clearly, she (weakly) prefers getting $\mathbf{1}$ for sure to playing r , and playing r to getting $\mathbf{0}$ for sure. So presumably, there is some number p between 0 and 1 such that she would as soon play r in the given game, as getting $\mathbf{1}$ with probability p , and $\mathbf{0}$ otherwise. This p is defined as her subjective probability for s .

This brings us to the difficulty. In strategic games, the play of the adversary depends on the strategies available to the protagonist. In most cases, the hypothetical strategy r of the protagonist is indeed hypothetical—it is not actually available to her. Making it available changes the game, and so may well change the likelihood that the adversary plays s . But if it is *not* available, how can the protagonist consider it?

To be sure, in games against nature, too, r may not be really available. But

²The protagonist is female. The other players are of indeterminate gender; we refer to them as “he,” to distinguish them from the protagonist.

there, nature is oblivious to the protagonist’s options; in the hypothetical situation created by adding r to the protagonist’s options, nature will play as in the original game. This is not so for strategic games³.

3. The Resolution

In this section, we state our result with a minimum of formality, while still striving for a maximum of transparency. Some compromises on both counts will be necessary. For a formal treatment, see Sections 4 and 5.

Following standard practice in decision theory, we rely on the primitive concept of “preference.” Preferences are applied to outcomes as well as to strategies (which, as noted above, correspond to Savage’s acts and A-A’s “horse lotteries”).

Before the difficulty can be resolved, we must clarify what it is that we require from a “resolution.” What, precisely, do we want subjective probabilities to do?

The first requirement is that the protagonist be able to use them to evaluate her options—in our terminology, her strategies. Specifically, starting out with the protagonist’s preferences,⁴ we want to construct utilities and subjective probabilities that “represent” those preferences: *inter alia*, such that she prefers one strategy to another if and only if its expected utility is greater.⁵ Call this the “representation” requirement (or property).

There is also another, equally important desideratum. In Section 2 we wrote that one may think of i ’s subjective probability for an event E as that number p such that i would as soon have a dollar with objective probability p as a dollar contingent on E . Differently put, the constructed subjective probabilities should be interchangeable, preference-wise, with objective probabilities that have the same numerical value. Call this the “interchangeability” requirement (or property).

Before proceeding, we introduce some terminology and notation. A *lottery* is an objective probability distribution.⁶ Unless otherwise indicated, a *strategy* is a strategy of the protagonist. A *consequence* is an outcome of the game. A *mixed strategy* is a lottery on strategies; a *mixed consequence* is a lottery on consequences.

³See Section 6.1 for further discussion of games against nature.

⁴Over her strategies and over the possible outcomes of the game.

⁵When the constructed subjective probabilities are assigned to the adversary’s choices, and the constructed utilities are assigned to the outcomes of the game.

⁶Anscombe and Aumann call this a *roulette lottery*; they also have *horse lotteries*, which are similar to what we here call a “strategy” of the protagonist. See Section 6.3.

We posit a preference order on the set Δ of all *hybrid* lotteries—those whose prizes may be *either* a pure strategy of the protagonist, *or* a consequence. This enables us to compare strategies with consequences, and so, to scale the utilities for strategies and for consequences to the same scale.

Operationally, such a lottery λ results either in (i) the outright selection of a specified pure outcome of the game, or in (ii) the game being played, with the protagonist choosing a specified pure strategy. More specifically, alternative (ii) results in awarding to the protagonist the outcome associated by the game with a specified strategy of the protagonist, combined with the strategy actually chosen by the adversary when playing the game. For each strategy s of the adversary, λ yields a mixed consequence λ_s in a natural way: if λ chose a consequence, then λ_s chooses the same consequence; and if λ chose a pure strategy r , then λ_s chooses the outcome of the game when the protagonist chooses r and the adversary chooses s . Note that all mixed consequences are in Δ , as are all mixed strategies; thus the preferences on Δ apply also to mixed consequences and to mixed strategies.

The following two assumptions are made:

N-M: The preference order satisfies the usual assumptions of von Neumann-Morgenstern utility theory; and

MONOTONICITY: If one lottery λ always yields a mixed consequence preferred to that yielded by another lottery λ' , no matter what the adversary does, then λ is preferred to λ' ; likewise for weak preference.⁷

We then have the following

MAIN THEOREM: There exists a function on the consequences, called a *utility* function, and a probability distribution on the adversary's strategies, called a *subjective probability distribution*, such that one lottery is preferred to another if and only if its expected utility is greater.

Here the expected utility is calculated using the objective probabilities that define the relevant lotteries, and the subjective probabilities for the adversary's strategies.

That the utilities and subjective probabilities in this theorem enjoy the representation property is immediate. They also enjoy the interchangeability property, since expectations do not change when subjective probabilities are replaced by objective ones.

⁷Preference or indifference.

4. Formal Treatment: Preliminaries

In the following, terms being defined are italicized.

The set of all probability distributions⁸ on a finite set A is denoted $\Delta(A)$. Note that if $\alpha, \alpha' \in \Delta(A)$ and $t \in (0, 1)$, then also $t\alpha + (1 - t)\alpha' \in \Delta(A)$. Abusing our notation, we write α and a interchangeably if α assigns probability 1 to a ; that is, we do not distinguish between a and a lottery that chooses a with certainty. No confusion should result.

A *preference order* \succsim on $\Delta(A)$ is a transitive⁹, reflexive¹⁰, and complete¹¹ binary relation on $\Delta(A)$. If $\alpha \succsim \beta$ and $\beta \succsim \alpha$, write $\alpha \sim \beta$ and say that α is *indifferent* to β . If $\alpha \succsim \beta$ and $\alpha \not\sim \beta$, write $\alpha \succ \beta$ and say that α is *preferred* to β . An *N-M utility* for \succsim is a real valued-function u on $\Delta(A)$ such that.

$$(4.1) \quad \alpha \succsim \alpha' \text{ iff } u(\alpha) \geq u(\alpha'), \text{ and}$$

$$(4.2) \quad u(t\alpha + (1 - t)\alpha') = tu(\alpha) + (1 - t)u(\alpha').$$

Various mutually equivalent axiom systems for N-M utility theory are available (von Neumann and Morgenstern [1944], Luce and Raiffa [1957], and others). We say that a preference order *satisfies the axioms of von Neumann-Morgenstern utility theory*—and call it an *N-M preference order*—if it satisfies any one of those systems.

PROPOSITION 4.3: An N-M preference order on $\Delta(A)$ has an N-M utility.

5. The Main Theorems

The viewpoint taken here is that of a single player, the *protagonist*, also called *Rowena*; it is her subjective probabilities for the strategy choices of the other players that we will define. Also, the preferences appearing below are hers, as are the utilities. It is convenient to combine all the other players into a single one, called *Colin*; we will see that no loss of generality is involved.

A *game* G consists of

a finite set R with members r (Rowena's *pure strategies*),

a finite set S with members s (Colin's *pure strategies*),

⁸Non-negative real-valued functions whose values sum to 1.

⁹ $\alpha \succsim \beta$ and $\beta \succsim \gamma$ imply $\alpha \succsim \gamma$.

¹⁰ $\alpha \succsim \alpha$.

¹¹ $\alpha \succsim \beta$ or $\beta \succsim \alpha$.

a finite set C with members c (*pure consequences*), and
a function $h : R \times S \rightarrow C$ (the *outcome function*¹²).

Members ρ, ρ', \dots of $\Delta(R)$ are called *mixed strategies of Rowena*; members γ, γ', \dots of $\Delta(C)$ are called *mixed consequences*; members λ, λ', \dots of $\Delta(R \cup C)$ (henceforth simply Δ) are called *hybrid lotteries* (or simply *lotteries*). If $\lambda \in \Delta$, set $\lambda = t\rho + (1-t)\gamma$, where $\rho \in \Delta(R)$, $\gamma \in \Delta(C)$, and $t \in [0, 1]$. For each pure strategy s of Colin, let $\rho_s \in \Delta(C)$ be the mixed consequence that results when Rowena plays ρ and Colin plays s , and let $\lambda_s := t\rho_s + (1-t)\gamma$ be the mixed consequence that results when Rowena uses the lottery λ and Colin plays s . Call a preference order \succsim on Δ *monotonic* if $\lambda \succsim \lambda'$ whenever $\lambda_s \succsim \lambda'_s$ for all s , and $\lambda \succ \lambda'$ whenever $\lambda_s \succ \lambda'_s$ for all s .

Now let $G = (R, S, C, h)$ be a game, \succsim a monotonic N-M preference order on Δ ; so in particular, $\succsim|_{\Delta(C)}$ is an N-M preference order on $\Delta(C)$, so has an N-M utility u , unique up to positive linear transformations.¹³

MAIN THEOREM: There exists a probability distribution p on S , such that for any hybrid lotteries λ, λ' ,

$$(5.1) \quad \lambda \succsim \lambda' \text{ if and only if } \sum_{s \in S} p_s u(\lambda_s) \geq \sum_{s \in S} p_s u(\lambda'_s).$$

Define $u_p(\lambda) := \sum_{s \in S} p_s u(\lambda_s)$ for all hybrid lotteries; then (5.1) becomes

$$(5.2) \quad \lambda \succsim \lambda' \text{ if and only if } u_p(\lambda) \geq u_p(\lambda').$$

In words, Rowena evaluates lotteries by their expected utility, when she ascribes probability p_s to Colin's choosing his pure strategy s .

6. Discussion

6.1. Games against Nature

In a game against nature, if we allow the protagonist to consider only a *specified* set of strategies, we get a theory formally identical to the one developed above. To be sure, in games against nature, she can always consider arbitrary “hypothetical” strategies, without running into the conceptual difficulty that motivates

¹²W.l.o.g. we could take $C = R \times S$ and let h be the identity; but nothing would be gained thereby, the notation would become more cumbersome, and the ideas less transparent.

¹³Multiplication by a positive constant and addition of an arbitrary constant.

¹⁴Clearly, this holds for one utility u iff it holds for all utilities u .

this paper (Section 2). Nevertheless, she might have difficulty in forming meaningful preferences between highly hypothetical options, and so prefer to restrict attention to the options actually available. Or, she might be reluctant to evaluate carefully acts that are clearly irrelevant.

Our framework differs from that of standard decision theory only in that the full set of acts (strategies) is replaced by a partial set. The main theorem says that as in the standard theory, we still get subjective probabilities and utilities that enjoy the representation and interchangeability properties.¹⁵

6.2. Non-Uniqueness

The subjective probabilities in the main theorem need not be unique; a simple example is a two-person game whose matrix has two identical columns, in which case the total subjective probability assigned to both columns can be divided between them in an arbitrary way. They are, however, *effectively* unique, in the sense that for a given utility function on the consequences, all subjective probability distributions satisfying the Main Theorem yield the same expected utility for each hybrid lottery. Formally, if p and p^* satisfy (5.1), then $u_p = u_{p^*}$. For the proof, see the end of Section 8.

6.3. The Decision Criterion

The main theorem provides an unequivocal answer to the question raised in the introduction: consistent decision making obeys the same basic logic in games against nature and in games of strategy, namely the logic of subjectively expected utility. This calls for (i) parallel assessment of the utilities of consequences and the probabilities of contingencies, both assessments being summarized in cardinal measures; (ii) the integration of utility and probability considerations through the calculation of expected utility.

The nature of preferences over strategies is brought out by the main theorem: one strategy r is “better” than another, r' , if and only if r carries higher expected utility, at the subjective probabilities that summarize the uncertainties of the situation. An important contribution of our analysis is the identification of the domain to which the principles of consistent decision making should be applied, namely *hybrid lotteries* over consequences and *strategies*.

¹⁵Equivalently, that a monotonic N-M preference order on the subset can be extended to a monotonic N-M preference order on the full set.

6.4. Observability and Operationalism

In the behaviorist tradition of Savage [1954, p.29] and of revealed preference theory [Samuelson, 1950], preferences are derived from choices: For the protagonist to prefer strategy ρ to strategy ρ' means that she chooses ρ when *confined to choosing* ρ or ρ' . This has the advantage of “operationalism” or “observability”—choices are in principle observable, preferences not. But in our context, it poses a serious problem: Confining her to choosing ρ or ρ' changes the game!

We have no unequivocal solution to this problem. In this work—as in most of decision theory—preferences among strategies are primitives; we have no totally satisfactory interpretation in terms of observable entities. Still, the following interpretative remarks provide some perspective.

While this is undoubtedly serious, it is not crippling. Certainly, operationalism and observability are vital principles. The question is, where do they fit in? Must the theory be *defined* in terms of observable entities, or must it have observable *implications*? Does one *start* with something observable, or *end* with it? One could perhaps go either way, but the second option—observable implications—certainly does not seem unreasonable.

Indeed, much of science operates in this way. None of the fundamental entities of physics—mass, energy, distance, gravitational force, etc.—are directly observable. These theoretical entities interact in complex ways, which eventually, sometimes, allows one to measure them. Here, analogously, one starts with the non-operational idea of preference among strategies; from this, one develops the theory of utility and subjective probability, which eventually enables an “operational” characterization of choices.

One could make another point. Suppose you must choose between x and y . You prefer x , and you choose x . Do you choose it because you prefer it, or do you prefer it because you choose it? Obviously the first makes more sense. Conceptually, preference comes before choice, it accounts for choice. So it makes sense to build the theory around preference and *derive* choice from it, just as actually happens in a person’s mind.

Whether or not one buys these arguments, preferences among strategies are in this work a primitive; they define the protagonist. The axioms (N-M and monotonicity) help her check their consistency and rationality. This is a mental process, similar to Savage’s (1954, pp. 28 and 20) “intermediate mode of interrogation” for normative application of his theory.

There is an alternative. Rather than taking preferences on hybrid lotteries as a primitive, one can take as primitives (i) an optimal set M in $\Delta(R)$, plus (ii)

a preference order on $M \cup \Delta(C)$. The interpretation is that operationally, the protagonist must choose one mixed strategy; the set of possible choices is M . A preference order on $M \cup \Delta(C)$ is also an operational concept. With appropriate assumptions, one can then derive subjective probabilities. Elsewhere, we plan to publish details, and compare such a treatment with that of this paper.

6.5. Reversal of Order

A-A use a formal assumption called “Reversal of Order in Compound Lotteries.” It says that if the outcome of a compound lottery depends in a *fixed, predetermined way* both on a roulette spin and a horse race, then it is immaterial whether the wheel is spun before or after the race. That is, a decision maker will be indifferent between these two possibilities, as long as he *cannot change* his horse bets after the roulette spin, or his roulette bets after the horse race. Though lacking substantive content in the A-A context, the assumption is required in their formalism¹⁶.

In our treatment, no separate assumption is necessary; it is already implicit in the formalism. Indeed, it follows from monotonicity on $\Delta(R \cup C)$ that if $\lambda_s = \lambda'_s$ for all s , then $\lambda \sim \lambda'$, which is an expression of the reversal of order assumption.

6.6. Value

The value to Rowena of the game G does not enter our analysis. It is natural to identify it with Rowena’s expected utility when she plays the game optimally, given how she thinks Colin will play; i.e., with $\max\{\sum_{s \in S} p_s u(r_s) : r \in R\}$. Unlike the minmax value of two-person zero-sum games, this does not call for Colin’s play to be “optimal” in any sense.

6.7. Resolutions that Don’t Work

6.7.1. Side Bets

A potentially operational approach to defining the protagonist’s subjective probabilities relies on “side bets.” In such a bet she gets, in *addition* to her payoff from the game, an amount δ if the adversary plays a specified strategy s ; nothing else is changed. One might then define her probability for s as that number p such that she would as soon opt for the side bet, as for δ with objective probability p .

¹⁶See Dreze [1987] for a context in which a similar assumption does have substantive content, namely the context of games against nature with moral hazard.

But it doesn't work; side bets may change the game. For example¹⁷, in the coordination game G_1 below, one may expect the Pareto dominant outcome BR . A side bet on L in the amount of $\delta = 8$ adds 8 to the row player's payoffs in column L , i.e., transforms the game to G_2 . If side bets "don't matter," we should expect BR in G_2 as well. By the same token, adding 8 to the column player's payoffs in row T of G_2 should not matter; this yields G_3 , commonly known as the "Stag Hunt"¹⁸. But here BR , which is Pareto dominated by TL , is far from compelling; indeed, Harsanyi and Selten [1987] select BR in G_1 and TL in G_3 .

	L	R		L	R		L	R
T	1, 1	0, 0		9, 1	0, 0		9, 9	0, 8
B	0, 0	7, 7		8, 0	7, 7		8, 0	7, 7
	G_1			G_2			G_3	

Side bets leave most equilibrium notions—including that of Nash (1951) and correlated equilibrium (Aumann, 1974)—invariant. Nevertheless, they subtly change incentives, as the example shows.

Some readers may hold that side bets do not change the game in any essential way; that in spite of appearances, G_3 would always be played like G_1 , and that that is common knowledge. In that case, of course, one *can* use side bets to define subjective probabilities. Put differently, one could restrict attention to players for whom it is common knowledge that G_3 is not essentially different from G_1 ; and again, with such a restriction, one can use side bets to define subjective probabilities.

6.7.2. Ignorance

In games against nature, the probability of a state s of nature is the number p such that Rowena is indifferent between a lottery yielding a dollar with objective probability p , and a bet r yielding a dollar if and only if s occurs. The fundamental difficulty set forth in Section 2 is that when s is a strategy of Colin in a strategic game, then r is usually not available to Rowena; making it available may change Colin's view of the game, and so Rowena's probability that he will choose s .

One might think it enough to consider a situation in which r really *is* available to Rowena, but *Colin does not know that it is*, so his choices—and Rowena's probabilities for his making those choices—will not be affected. But that is not

¹⁷Communicated by S. Hart. To avoid difficulties, assume dollar payoffs and linear utilities.

¹⁸See O'Neill [1994, pp.1004-5] for a discussion of this game and some of the literature on it.

very satisfactory. What *does* Colin know? Does he consider it possible that r is available to Rowena? If he does, then already that changes the game, and we have the same difficulty as before. If he considers it impossible—*knows* that r is unavailable—how can it be available? There is a basic incoherence in situations where something false is “known.”

7. Affine Monotonic Functions

For points x, y in \mathbb{R}^n , write $x \gg y$ if $x_i > y_i$ for all i , and write $x \geq y$ if $x_i \geq y_i$ for all i . A real-valued function f from a convex set D in \mathbb{R}^n to \mathbb{R} is called *affine* if $f(tx + (1-t)y) = tf(x) + (1-t)f(y)$ for all x, y in D and t in $(0, 1)$. It is called *monotonic* if $x \gg y$ implies $f(x) > f(y)$, and $x \geq y$ implies $f(x) \geq f(y)$, for all x, y in D .

PROPOSITION 7.1: Let H be a convex subset of \mathbb{R}^n , and f an affine monotonic real-valued function on H . Assume¹⁹ that there are points z and z' in H with $z \gg z'$. Then there exist non-negative q_1, \dots, q_n , not all of which vanish, and a real q_0 , such that $f(x) = q_0 + \sum_{i=1}^n q_i x_i$ for all x in H .

In this section we prove 7.1; readers willing to accept the proposition on faith may proceed at once to the proof of the main result in the next section.

The origin $(0, \dots, 0)$ of \mathbb{R}^n is denoted $\mathbf{0}$. A *linear subspace* M of \mathbb{R}^n is a subset of \mathbb{R}^n that, together with any two points x, y in it, and any real number t , contains $x+y$ and tx . A function f on M is *linear* if $f(x+y) = f(x)+f(y)$ and $f(tx) = tf(x)$ for all x, y in M and all real t . Note that a function on M is linear if and only if it is affine. A *hyperplane* in M is a set of the form $\{x \in M : f(x) = t\}$, where f is a linear function on M that does not vanish identically, and t is a constant.

LEMMA 7.2: If D is an open convex set in \mathbb{R}^n and M a linear subspace of \mathbb{R}^n that does not meet D , then there is a hyperplane in \mathbb{R}^n that includes M and does not meet D ; i.e., a linear function on \mathbb{R}^n that vanishes on M and is positive on D .

PROOF: Eggleston [1958], p.19, Theorem 7.

Define $\mathbb{R}_{++}^n := \{x \in \mathbb{R}^n : x \gg \mathbf{0}\}$.

LEMMA 7.3: Let L be a linear subspace of \mathbb{R}^n that intersects \mathbb{R}_{++}^n . Then any monotonic linear function f on L may be extended to a monotonic linear function on all of \mathbb{R}^n .

¹⁹The proposition is true as it stands even without this assumption, but the proof lies deeper.

PROOF: Define $L_0 := \{x \in L : f(x) = 0\}$. Because f is monotonic, L_0 cannot meet \mathbb{R}_{++}^n . Use 7.2 to find a linear function f^* on \mathbb{R}^n that vanishes on L_0 and is positive on \mathbb{R}_{++}^n . Let x_1 be a point in $\mathbb{R}_{++}^n \cap L$. Then $f(x_1) > 0$. Similarly $f^*(x_1) > 0$. Possibly redefining f^* by multiplication by a positive constant, we may take $f(x_1) = f^*(x_1)$. So f and f^* coincide on L_0 and on x_1 , which together span L ; so they coincide on L . Thus f^* extends f , and since it is positive on \mathbb{R}_{++}^n , it is monotonic. This completes the proof of 7.3.

PROOF OF PROPOSITION 7.1. By possibly applying a translation, we may suppose w.l.o.g. that $\mathbf{0}$ is in the interior of H , relative to the smallest linear manifold L that includes H . Since $\mathbf{0} \in L$, it follows that L is a linear space, and there is a unique extension of f to a linear function f' on L . Both z and z' are in H and so in L , so $z - z'$ is in L , so L intersects \mathbb{R}_{++}^n , since $z - z' \gg \mathbf{0}$. To see that f' is monotonic, let x in L be $\gg \mathbf{0}$. Then for sufficiently small positive ε , we have $\varepsilon x \in L$ and $\varepsilon x \gg \mathbf{0}$, so $\varepsilon f'(x) = f'(\varepsilon x) = f(\varepsilon x) > 0$, so $f'(x) > 0$. Similarly, if x in L is $\geq \mathbf{0}$, then $f'(x) \geq 0$. Thus f' is monotonic as well as linear. Applying 7.3, we obtain an extension f'' of f' from L to a linear monotonic function on all of \mathbb{R}^n . Let $f''(x) = \sum_{i=1}^n q_i x_i$ for all x in \mathbb{R}^n ; all linear functions on \mathbb{R}^n have this form. It cannot be that all the q_i vanish, for then $f''(1, \dots, 1) = 0 = f(\mathbf{0})$, contrary to the monotonicity of f'' . To show that the q_i are non-negative, suppose, say, that $q_1 < 0$. Then $f''(1, 0, \dots, 0) = q_1 < 0$, again contradicting the monotonicity of f'' . So the q_i are non-negative and do not all vanish, proving²⁰ the proposition.

8. Proof of the Main Theorem

By Proposition 4.3, \succsim has an N-M utility u ; from 4.1, it follows that

$$(8.1) \quad \lambda \succ \lambda' \text{ iff } u(\lambda) > u(\lambda').$$

W.l.o.g., we assume that there are pure consequences d and d' with²¹

$$(8.2) \quad d \succ d'.$$

Set $S := \{s_1, \dots, s_n\}$. With each lottery λ in Δ , associate the point $u(\lambda_S) := (u(\lambda_{s_1}), \dots, u(\lambda_{s_n}))$ in \mathbb{R}^n . Note that when λ chooses some mixed consequence γ for sure (rather than playing the game with positive probability)—i.e., when $\lambda \in \Delta(C)$ —then $\lambda_{s_1} = \dots = \lambda_{s_n} = \gamma$, so

$$(8.3) \quad u(\lambda_S) := (u(\lambda_{s_1}), \dots, u(\lambda_{s_n})) = (u(\gamma), \dots, u(\gamma)).$$

²⁰The term q_0 is due to the translation at the beginning of the proof.

²¹If all consequences are indifferent, take u to be identically 0, and p an arbitrary distribution.

Let H be the set of all the points $u(\lambda_S)$ when λ ranges over Δ . By 4.2, we have $tu(\lambda_s) + (1-t)u(\lambda'_s) = u(t\lambda_s + (1-t)\lambda'_s) = u((t\lambda + (1-t)\lambda')_s)$ for each s in S and t in $[0, 1]$ so

$$(8.4) \quad tu(\lambda_S) + (1-t)u(\lambda'_S) = u((t\lambda + (1-t)\lambda')_S),$$

so H is convex. Moreover, if d and d' are as in 8.2, then by 8.3, H contains the points $z := (u(d), \dots, u(d))$ and $z' := (u(d'), \dots, u(d'))$, so by 4.1, $z > z'$. Thus H satisfies the hypotheses of 7.1.

Now define a function f on H by $f(x) := u(\lambda)$ for any λ for which $u(\lambda_S) = x$; that there is such a λ follows from $x \in H$, and that it doesn't matter which one we use follows from monotonicity: If $u(\lambda_S) = u(\lambda'_S)$, then $\lambda \sim \lambda'$, by 5.1 and monotonicity, so $u(\lambda) = u(\lambda')$, by 8.1. Next, we show that f is affine and monotonic. Indeed, let $x = u(\lambda_S)$ and $x' = u(\lambda'_S)$; then by 8.4 and 4.2,

$$\begin{aligned} f(tx + (1-t)x') &= f(tu(\lambda_S) + (1-t)u(\lambda'_S)) = f(u((t\lambda + (1-t)\lambda')_S)) \\ &= u(t\lambda + (1-t)\lambda') = tu(\lambda) + (1-t)u(\lambda') = tf(x) + (1-t)f(x'), \end{aligned}$$

$$(8.5) \quad f \text{ is affine.}$$

To show that

$$(8.6) \quad f \text{ is monotonic,}$$

first let $x \geq x'$; then $u(\lambda_S) \geq u(\lambda'_S)$, so $u(\lambda_s) \geq u(\lambda'_s)$ for all s , so $\lambda_s \succsim \lambda'_s$ for all s , (by 4.1), so $\lambda \succsim \lambda'$ (by monotonicity), so $f(x) = u(\lambda) \geq u(\lambda') = f(x')$, by 8.1. Similarly, if $x \gg x'$, then $f(x) > f(x')$, proving 8.6. So by 7.1, there exist non-negative q_1, \dots, q_n , not all of which vanish, and a real q_0 , such that $f(x) = q_0 + \sum_{i=1}^n q_i x_i$ for all x in H . So $u(\lambda) = q_0 + \sum_{i=1}^n q_i u(\lambda_S)_i = q_0 + \sum_{i=1}^n q_i u(\lambda_{s_i})$. So by 8.1,

$$(8.7) \quad \lambda \succsim \lambda' \text{ if and only if } \sum_{i=1}^n q_i u(\lambda_{s_i}) \geq \sum_{i=1}^n q_i u(\lambda'_{s_i}).$$

Set $p_i = q_i / \sum_{i=1}^n q_i$; the denominator does not vanish because the q_i are non-negative and do not all vanish. So the p_i are non-negative and sum to 1; that is, they constitute a probability distribution on $S = \{s_1, \dots, s_n\}$. The conclusion of the Main Theorem then follows from 8.7.

To prove the effective uniqueness (Section 6.2), let λ be a hybrid lottery. Preference-wise, λ must be between the most preferred and the least preferred consequence; so there is a mixed consequence ρ with $\lambda \sim \rho$. then 5.2 yields $u_p(\lambda) = u_p(\rho) = u(\rho) = u_{p^*}(\rho) = u_{p^*}(\lambda)$, as was to be proved.

9. Literature

Luce and Raiffa [1957, p.306] were among the earliest²² to suggest assigning subjective probabilities to an adversary’s choices in a strategic game; they wrote as follows: “The problem of individual decision making under uncertainty can be considered a one-person game against a neutral nature. Some of these ideas can be applied indirectly to individual decision making ... where the adversary is not neutral but a true adversary. ... One *modus operandi* for the decision maker is to generate an *a priori* probability distribution over the ... pure strategies ... of his adversary by taking into account both the strategic aspects of the game and ... ‘psychological’ information ... about his adversary, and to choose an act which is best against this ... distribution.” They go on to explore the idea of “side bets” (see 6.6.1 above), noting some difficulties with it, and informally suggesting a possible way around them. No formal model was developed, and no definite conclusion reached.

It appears that Armbruster and Böge [1979] and Böge and Eisele [1979] were the first to construct formal models in which each player directly²³ assigns subjective probabilities to the strategy choices of the others. A relatively early application of this idea is Brandenburger and Dekel [1987]. The representation of the value of a game to a player as a subjectively expected utility is implicit in the work of Nau and McCardle [1990].

Kadane and Larkey [1982] claim that the problem of a player in a game is no different from any other one-person decision problem. In particular, they suggest abandoning altogether all notions of equilibrium and any “interactive” analysis (i.e., analysis of each player’s beliefs about the others’ beliefs about his actions and beliefs). Instead, they propose simply that each player form, in some unspecified and unrestricted way, a probability distribution over the other players’ strategies, and then maximize against that. To form the probabilities, they suggest using disciplines like cognitive psychology rather than decision or game theory.

In the precisely opposite direction, Mariotti [1995, p.1108] writes that “a di-

²²They cite an earlier paper by Hodges and Lehmann [1952] who suggest that a player in a two-person zero-sum game might assign subjective probabilities to the eventuality that his adversary will make a “mistake.” But this is not really in the spirit of this paper, nor of Luce and Raiffa’s suggestion.

²³Previously, Aumann [1974] had already used subjective probability in analyzing games; but in that analysis, players use “subjectively mixed strategies”—peg their pure strategy choices on events (like outcomes of horse races) whose probability is not agreed upon—rather than simply assigning a subjective probability to the other players’ choices.

voice is required between game theory and individual decision theory ... strategic decision principles may be radically different from individual decision-theoretic principles.”

Some of the ideas underlying the current work appear already in Mariotti’s excellent paper. Inter alia, that a strategy in a game corresponds to an act in Savage’s one-person decision theory²⁴; that “only some acts (strategies) are feasible for each player in a given game,” and the players should “rank only the strategies available in that game” (p.1102); and the difficulty discussed at 6.4 above, at which he hints in his discussion (p.1108(c)).

None of the authors cited above, nor indeed others, seem aware of the fundamental difficulty addressed by this paper: that the standard construction of subjective probabilities does not apply to strategic games.

10. References

Anscombe, F. J., and R. J. Aumann, (1963), “A Definition of Subjective Probability,” *Ann. Math. Stat.* 34, 199-205 (**A-A** in the text).

Armbruster, W. and W. Boege, (1979), “Bayesian Game Theory,” in *Game Theory and Related Topics*, edited by O. Moeschlin and D. Pallaschke, Amsterdam: North Holland.

Aumann, R. J. (1974), “Subjectivity and Correlation in Randomized Strategies,” *J. Math. Econ.* 1, 67-96.

Boege, W. and T. Eisele, (1979), “On Solutions of Bayesian Games,” *International Journal of Game Theory* 8, 193-215.

Brandenburger, A., and E. Dekel (1987), “Rationalizability and Correlated Equilibria,” *Econometrica* 55, 1391-1402.

Dreze, J. H. (1961), “Les fondements logiques de l’utilité cardinale et de la probabilité subjective,” in *La Décision*, Paris: CNRS, 73-87.

——— (1987), *Essays on Economic Decisions under Uncertainty*, Cambridge, at the University Press.

Eggleston, H. G. (1958), *Convexity*, Cambridge, at the University Press.

de Finetti, B. (1937), “La Prevision: Ses Lois Logiques, ses Sources Subjectives,” *Ann. Inst. H. Poincaré* 7, 1-68.

²⁴And so also to a horse lottery in A-A.

- Harsanyi, J. C., and R. Selten (1987), *A General Theory of Equilibrium Selection in Games*, Cambridge, Mass.: MIT Press.
- Kadane, J. B., and P. D. Larkey (1982), "Subjective Probability and the Theory of Games," *Manag. Sci.* 28, 113-120.
- Luce, R. D., and H. Raiffa (1957), *Games and Decisions*, New York: John Wiley.
- Hodges, J. L. Jr., and E. L. Lehmann (1952), "The Uses of Previous Experience in Reaching Statistical Decisions," *Ann. Math. Stat.* 23, 396-407.
- Mariotti, M. (1995), "Is Bayesian Rationality Compatible with Strategic Rationality?" *Econ. J.* 105, 1099-1109.
- Nash, J. F. (1951), "Non-cooperative Games," *Ann. Math.* 54, 286-295.
- Nau, R. F. and K. F. McCardle (1990), "Coherent Behaviour in Non-cooperative Games", *J. Econ. Th.* 50, 424-444.
- von Neumann, J., and O. Morgenstern (1944), *Theory of Games and Economic Behavior*, Princeton: Princeton University Press.
- O'Neill, B., (1994), "Game Theory Models of Peace and War," in *Handbook of Game Theory with Economic Applications, Vol.2*, R.J. Aumann and S.Hart, eds., Amsterdam: Elsevier, 996-1053.
- Ramsey, F. P., (1931), "Truth and Probability," Chapter VII of *The Foundations of Mathematics*, edited by R.B. Braithwaite, London: Routledge and Kegan Paul, 156-198.
- Samuelson, P. A., (1950) "The Problem of Integrability in Utility Theory," *Economica* NS 17, 355-385.
- Savage, L. J. (1954), *The Foundations of Statistics*, New York: John Wiley.

Département des Sciences Économiques
de l'Université catholique de Louvain
Institut de Recherches Économiques et Sociales

Place Montesquieu, 3
1348 Louvain-la-Neuve, Belgique