

A Derivation of Expected Utility Maximization in the Context of a Game*

Itzhak Gilboa[†] and David Schmeidler[‡]

Revised – April 2002

Abstract

A decision maker faces a decision problem, or a game against nature. For each probability distribution over the state of the world (nature's strategies), she has a weak order over her acts (pure strategies). We formulate conditions on these weak orders guaranteeing that they can be jointly represented by expected utility maximization with respect to an almost-unique state-dependent utility, that is, a matrix assigning real numbers to act-state pairs. As opposed to a utility function that is derived in another context, the utility matrix derived in the game will incorporate all psychological or sociological determinants of well-being that result from the very fact that the outcomes are obtained in a given game.

JEL classification numbers: C70, D80, D81

1 Introduction

1.1 Motivation

Do players maximize expected utility when playing a game? The experimental outcomes involving ultimatum and dictator games might seem to suggest

*We thank an associate editor and two anonymous referees for comments and references. This work was partially supported by ISF Grant No. 790/00.

[†]Tel-Aviv University and Cowles Foundation, Yale University. igilboa@post.tau.ac.il

[‡]Tel-Aviv University and The Ohio State University. schmeid@tau.ac.il

that they do not. (See Guth and Tietz (1990) and Roth (1992) for surveys.) For instance, a player who moves second in an ultimatum game, and rejects an offer of a positive amount of money, evidently does not maximize her monetary payoff. Similarly, a dictator in a dictator game, who chooses to leave some money to her dummy opponent, fails to maximize her payoff under conditions of certainty, let alone her expected payoff under conditions of uncertainty.

Some authors argue that these experimental results constitute a violation of game theoretic predictions. Indeed, if one insists that the utility function be defined over monetary payoffs alone, such a conclusion appears unavoidable. But many game theorists hold that the utility function need not be defined on monetary prizes alone. Indeed, an “outcome” should specify all the relevant features of the situation, including feelings of envy, guilt, preferences for fairness, and so forth. Moreover, recent developments in economic theory call for explicit modeling of such determinants of utility. (See, for instance, Frank (1988), Elster (1998), Rabin (1998), and Loewenstein (2000).) Further, if one adopts a purely behavioral approach, one has no choice but to incorporate into the utility function all psychological and sociological effects on well-being. The very fact that, say, a dictator prefers taking less money to taking more money implies that the utility of the former exceeds that of the latter. As long as players do not violate the axioms of von-Neumann and Morgenstern (vNM, 1944), they can be described as if they are maximizing the expected value of an appropriately chosen utility function. From this viewpoint, the experimental results of dictator and ultimatum games might challenge the implicit assumption that monetary payoff is the sole determinant of utility, but not the assumption of expected utility maximization itself.

We find that this argument is essentially correct: the debate aroused by dictator and ultimatum games is about determinants of the utility function, not about expected utility theory (EUT). Yet, we do not believe that vNM’s

axiomatic derivation of EUT is a very compelling argument in this context. vNM's result assumes a preference relation over lotteries with given probabilities, and derives a utility function over outcomes, such that the maximization of its expectation represents preferences over lotteries. vNM then assumed that, when players evaluate mixed strategies in a game, they use the same utility function for the calculation of expected payoff, and attempt to maximize this expectation. Thus, the vNM derivation implicitly assumes that the utility function that one obtains in the context of a single person decision problem will apply to the context of a game.

This assumption seems implausible precisely in the context of games such as ultimatum and dictator, where utility is heavily dependent on interpersonal comparisons and interactions. For instance, if player two considers an outcome of 10% of the pie, she cannot ignore the fact that player one (the divider) is about to pocket 9 times as much. Similarly, player one (the dictator) in a dictator game cannot be assumed to treat the outcome "I get \$90" as equivalent to "I chose to take \$90 and to leave \$10 to my opponent." Preferences over fairness distinguish the former from the latter. Moreover, the very fact that the dictator has *chosen* a particular division of the money implies that she might experience guilt even if she has no preference for fairness *per se*. Finally, suppose that player two in an ultimatum game chooses to reject an offer not because it is unfair, but because she finds it insulting. That is, she does not envy player one, but she finds that he should be punished for his greed. In this case, she distinguishes between "I get \$10, player one gets \$90, and this was decided by Nature" and "I get \$10, player one gets \$90, and this was decided by Player one." Such distinctions are precisely about the difference between a single-player decision problem and a game. If one were to measure a player's utility over such outcomes in a laboratory, one would have to generate outcomes that simulate all the interactive effects of a game. That is, one would have to measure utility *in the context of the game itself*.

Similar issues arise when a single player is concerned. Consider, for instance, the effect of regret. It has long been argued that regret may color the way individuals evaluate outcomes. (See, for instance, Luce and Raiffa (1957), Loomes and Sugden (1982), and Gul (1991).) Thus, the utility function of a certain outcome, when measured in isolation, may not reflect the way this outcome is perceived in a game. “Getting \$10” is not the same as “Getting \$10 when I could have gotten \$20.” In order to measure the relevant utility of the latter, one would have to simulate the entire choice situation, that is, to measure utility in the context of the game.

In order to defend the expected utility paradigm in face of experimental evidence as well as of the theoretical considerations mentioned above, it does not suffice to show that it *can* explain the data with an appropriate definition of the utility function. One needs to show that this new definition also relies on sound axiomatic foundations. That is, one needs an axiomatic derivation of EUT that would parallel that of vNM, but will only use preferences in the game itself as data.

1.2 The present contribution

In this paper we axiomatize expected utility maximization in the context of a game.¹ We assume that every player can rank his pure strategies, given *any* distribution over the pure strategies of the other players. Such a distribution is interpreted as the player’s subjective beliefs. In a two-person game, these beliefs might also coincide with the other player’s mixed strategy. Equivalently, one may consider a single person decision problem under uncertainty (a “game against nature”), where, for each vector of probabilities over the state of nature (representing the decision maker’s beliefs), the decision maker has a weak order over the possible acts. The set of acts may be finite or infinite, and it is not assumed to have any algebraic, topological, or other structure. The set of states of nature may also be finite or infinite.

¹See Hammond (1997) for a different approach, employing classes of games.

Pairs of acts and states (or combinations of pure strategies) can be thought of as defining outcomes. In the formal model we do not assume any specification of physical outcomes. When such specifications exist, they may be part of the definition of outcomes, but they need not summarize all the utility-relevant information pertaining to the outcomes. In particular, one may not infer dominance relations between strategies from physical descriptions of outcomes.

When a player compares two acts, given a distribution over the states of nature, she may be viewed as comparing lotteries, namely, distributions over outcomes. Observe, however, that we do not assume that the player can compare any pair of lotteries defined over the possible outcomes. First, we do not consider lotteries that assign positive probability to outcomes that result from different acts. Second, we do not assume that the player can compare any two such lotteries. Rather, we assume that the player can compare only lotteries that induce the same marginal distribution over states. In particular, the data assumed in our results will not include a comparison of a certain outcome (i.e., a degenerate lottery) to a non-degenerate lottery.

We assume that the rankings over the acts (the player's pure strategies) satisfy two axioms that relate preferences given different beliefs (different mixed strategies of the opponent): first, we assume *convexity*: if act a is preferred to b given probability p , as well as given probability q , then the same preference will be observed for any convex combination of p and q . Second, we assume *continuity*: if a is strictly preferred to b given belief p , the same preference should hold in a neighborhood of p . Finally, we also need an axiom of *diversity*, requiring that any four pure acts can be ranked, in any given strict order, for at least one belief vector p . (See the following section for a more precise formulation of the axioms.) These axioms imply that there is a utility matrix, such that, for every belief (opponent's mixed strategy) p , the decision maker (player) ranks her acts (pure strategies) according to their expected utility, computed for the relevant p .

We do not claim that the expected utility paradigm is broad enough to encompass all types of psychological or social payoffs. Indeed, there are modes of behavior that would violate expected utility maximization for *any* utility function, and would therefore also violate our axioms. Our goal is not to argue for the universality of the expected utility paradigm, but to contribute to the precise delineation of its scope of applicability.

Since in our model the utility we derive is defined over act-state pairs, it is a state-dependent utility function. (See Dreze (1961) and Karni, Schmeidler, and Vind (1983).) This also implies that there is some freedom in the choice of the utility function: one may add a separate constant to each column in the matrix without changing the expected utility rankings. Indeed, our utility matrix is unique up to such shifts, and up to multiplication of the entire matrix by a positive number.

The diversity axiom implies that the matrix we obtain satisfies a certain condition, which we dub “diversification”: no row in the matrix is dominated by an affine combination of (up to) three other rows in it. In particular, it does not allow domination relations between pure strategies. However, in the absence of the diversity assumption, the other axioms do not imply the existence of the numerical representation we seek.

Further discussion of our axioms and the result is deferred to Section 3. We now turn to the formal statement of the model and result.

2 Result

The result presented in this section is reminiscent of the main results in Gilboa and Schmeidler (1997, 1999). All these results derive a representation of a family of weak orders by a matrix of real numbers, as follows. The objects to be ranked corresponds to rows in the matrix. A “context”, which induces a weak order over these objects, is defined by a function, attaching a real number to each column. Given such a context, the ranking corresponding

to it is represented by the inner products of the context with each of the rows in the matrix. While our new results are similar in spirit to those in previous papers, some differences exist. In particular, the extension to an infinite state space is new.

There are several reasons for which one may be interested in an infinite space. First, there are many situations in which Nature or other players actually have infinitely many strategies. Second, even when other players have finitely many pure strategies, one need have preferences that are linear in the probabilities defining their mixed strategies. In such situations, we may introduce mixed strategies explicitly as columns, and still assume that the player's preferences are linear in her own beliefs over these columns. However, such a construction naturally leads to an infinite matrix. Finally, a player may have psychological payoffs depending on the intentions, beliefs, and more generally, type of other players. These may also call for infinitely many columns in the matrix.²

Assume that a decision maker is facing a decision problem with a non-empty set of acts A and a measurable space of states of the world (Ω, Σ) , where Σ is a σ -algebra of subsets of Ω . Further, assume that Σ includes all singletons. Let $\mathbf{B}(\Omega, \Sigma)$ be the space of bounded Σ -measurable real-valued functions on Ω . Recall that $\mathbf{ba}(\Omega, \Sigma)$, the space of finitely additive bounded measures on Σ , is the dual of $\mathbf{B}(\Omega, \Sigma)$. Let \mathbf{P} denote the subset of $\mathbf{ba}(\Omega, \Sigma)$ consisting of finitely additive probability measures on Σ . Assume that, for every probability measure $p \in \mathbf{P}$, the decision maker has a binary preference relation $\%_p$ over A . As usual, we denote by \succ_p and \sim_p the asymmetric and symmetric parts of $\%_p$, respectively. The axioms on $\{\%_p\}_{p \in \mathbf{P}}$ are:

A1 Ranking: For every $p \in \mathbf{P}$, $\%_p$ is complete and transitive on A .

A2 Combination: For every $p, q \in \mathbf{P}$ and every $a, b \in A$, if $a \%_p b$ ($a \succ_p b$) and $a \%_q b$, then $a \%_{\alpha p + (1-\alpha)q} b$ ($a \succ_{\alpha p + (1-\alpha)q} b$) for every $\alpha \in (0, 1)$.

A3 Continuity: For every $a, b \in A$ the set $\{p \in \mathbf{P} \mid a \succ_p b\}$ is open in

²We thank the associate editor for this last point.

the relative weak* topology.

A4 Diversity: For every list (a, b, c, d) of distinct elements of A there exists $p \in \mathbf{P}$ such that $a \succ_p b \succ_p c \succ_p d$. If $|A| < 4$, then for any strict ordering of the elements of A there exists $p \in \mathbf{P}$ such that \succ_p is that ordering.

Axioms 1 and 3 are rather standard. Axiom 2 states that the set of beliefs, for which act a is preferred over act b , is convex. It is rather natural if payoffs, including psychological and sociological ones, do not depend on probability mixtures. By contrast, it should be expected to fail if probability mixing itself generates affective reactions, or if it defines social characteristics. The main justification for Axiom 4 is mathematical (see below). Observe that it precludes cases in which one act dominates another.

We need the following definition: a matrix of real numbers is called *diversified* if no row in it is dominated by an affine combination of three (or less) other rows in it. Formally:

Definition: A matrix $u : A \times \Omega \rightarrow \mathbf{R}$, where $|A| \geq 4$, is *diversified* if there are no distinct four elements $a, b, c, d \in A$ and $\lambda, \mu, \theta \in \mathbf{R}$ with $\lambda + \mu + \theta = 1$ such that $u(a, \cdot) \leq \lambda u(b, \cdot) + \mu u(c, \cdot) + \theta u(d, \cdot)$. If $|A| < 4$, u is diversified if no row in u is dominated by an affine combination of the others.

We can now state

Theorem 1 : *The following two statements are equivalent :*

- (i) $\{ \%_p \}_{p \in \mathbf{P}}$ satisfy A1, A2, A3, A4;
- (ii) *There exists a diversified matrix $u : A \times \Omega \rightarrow \mathbf{R}$ such that $u(a, \cdot) \in \mathbf{B}(\Omega, \Sigma)$ for all $a \in A$ and*

$$\text{for every } p \in \mathbf{P} \text{ and every } a, b \in A,$$

$$(**) \quad a \%_p b \quad \text{iff} \quad \int_{\Omega} u(a, \cdot) dp \geq \int_{\Omega} u(b, \cdot) dp ,$$

Furthermore, if (i) (equivalently, (ii)) holds, the matrix $u(\cdot, \cdot)$ is unique in the following sense: a matrix $w : A \times \Omega \rightarrow \mathbf{R}$ with $w(a, \cdot) \in \mathbf{B}(\Omega, \Sigma)$ for all

$a \in A$ satisfies $(**)$ iff there are a scalar $\lambda > 0$ and a function $v \in \mathbf{B}(\Omega, \Sigma)$ such that $w(a, \cdot) = \lambda u(a, \cdot) + v$ for all $a \in A$.

The proof of this theorem is given in an appendix.

We do not know of a set of axioms that are necessary and sufficient for a representation as in $(*)$ by a matrix u that need not be diversified. We do know that dropping A4 will not do. (The counter-examples in Gilboa and Schmeidler (1997, 1999) can be easily adapted to our case.) It will be clear from the proof that weaker versions of A4 suffice for a representation as in $(*)$. Ashkenazi and Lehrer (2001) also offer a condition that is weaker than A4, and that also suffices for a similar representation. The diversity axiom is stated here in its simplest and most elegant form, rather than in its mathematically weakest form.

3 Discussion

Fishburn (1976) and Fishburn and Roberts (1978) provide derivations of expected utility maximization in the context of a game. In these papers, a player is assumed to have preferences over lotteries that are generated by her own mixed strategies and by mixed strategies of the opponents. These results do not suffice for our purposes for two reasons. First, they assume that all lotteries, obtained by independent mixed strategies, can be compared. But each player in a game can only choose her own strategies. Thus, to make such preferences observable one would, again, have to resort to experimental settings that are external to the game. Second, a player's preference over her own mixed strategies has been criticized as shaky data. It is not clear when a player's actual choices (of pure strategies) reflect preferences over mixed strategies. Moreover, it has been argued that players never actually play mixed strategies. (See Rubinstein (2000).) We do not assume preferences

over a player's own mixed strategies,³ and interpret a mixed strategy of a player merely as the beliefs of other players regarding her (pure strategy) choice. (See Aumann and Brandenburger (1995).)

The question of observability

We assume that a player can rank pure strategies given *any* belief over the opponents' strategies. There are situations in which this assumption is rather plausible. For instance, suppose that a player is anonymously matched with other players and is given statistical data regarding past plays of the game by other players from the same population. In an experimental set-up, one may induce any probability vector p as the player's beliefs. Observe that this assumption is compatible with the player's belief in rationalizability: as dictator and ultimatum games indicate, knowledge of the physical outcomes of the game does not imply knowledge of dominance relations for other players.

By contrast, consider a situation in which a player is matched with another player, whom she knows. There might be beliefs p that the player would not entertain regarding her opponent. For instance, if Mary is happily married to John, and the two are matched to play an ultimatum game, Mary might be convinced that John, as player 1, will never make an ungenerous offer. In this case, we will never know what she would do if she did believe that John is ungenerous. In particular, we will not be able to tell whether Mary is nice to John because he is generous in his dealings with her, or because she will like him even if he treats her badly.

To address this difficulty, we point out that our result can be extended. Our theorem holds also when beliefs are restricted to a convex subset of probability measures. In this case the continuity axiom implies that the set of probability measures has a nonempty interior relative to \mathcal{P} . Further, one may strengthen the result by weakening the continuity axiom so that it applies only relative to the convex set of probability measures. Observe,

³Naturally, we also do not derive representation of preferences of a player over her own mixed strategies.

however, that requiring the diversity axiom on a subset of P is more restrictive than requiring it on P in its entirety. Moreover, if one insists that preferences can only be elicited given the players' actual beliefs, our approach does not apply.⁴

One may wonder how can an experimenter observe a binary relation, whereas, in reality, only a choice out of a given set is made by the player.⁵ One way to experimentally observe an entire binary relation is the following. A player is told that not all acts need be available, and asked to provide a ranking, such that the best available act according to this ranking will eventually be played. If the additional complication in the experimental procedure does not introduce new affective reactions, the player will reveal a meaningful binary relation.

Roles of axiomatizations

Axiomatic derivations can be useful also when the presumed data are not directly or easily observable. First, one may use an axiomatic system such as ours for normative purposes. In this case one need not verify that the axioms are satisfied by a given player, or even observe the player's preferences at all. Rather, a player who finds our axioms compelling will be glad to know that they are consistent and, furthermore, that the only way to satisfy them is by maximization of expected utility, for an appropriately defined utility function. The player may then map her utility function, incorporating psychological and social factors, and use it for decision making.

Second, the rhetorical use of axioms can also convince one that expected utility theory is a useful descriptive tool even when it cannot be directly tested. Specifically, some researchers cast doubts on the usefulness of the expected utility paradigm in the presence of psychological and social factors. Assume that such a researcher finds that these factors do not lead to violations of our axioms. In this case, she may conclude that maximization of

⁴Hammond (1997) addresses this problem. His approach, however, requires choices in hypothetical games, and not just in the game that is actually being played.

⁵We thank an anonymous referee for raising this point.

expected utility can still be a valid description of behavior even when such factors are present.

Third, an axiomatic derivation delineates the scope of refutability of a theory. Thus, should one suspect that any mode of behavior can be justified as expected utility maximization for a carefully chosen utility function, axioms such as ours show that this is not the case. More generally, our axioms show precisely what is meant by expected utility maximization when the utility function may incorporate various game-specific determinants of utility.

Finally, one may use axioms to define and measure a utility function. The definition of the utility function serves a theoretical purpose, and characterizes the degree to which this function is unique. The measurement of utility is done by elicitation of observable preference data. This process can be greatly simplified using the structure provided by our result.

Appendix: Proof

Proof of Theorem 1:

Theorem 1 is reminiscent of the main results in Gilboa and Schmeidler (1997, 1999; see also 2001). Although the spaces discussed are different, some steps in the proof are practically identical. We therefore provide here only a sketch of steps that appear elsewhere in detail.

We present the proof for the case $|A| \geq 4$. The proofs for the cases $|A| = 2$ and $|A| = 3$ will be described as by-products along the way. For $u \in \mathbf{B}(\Omega, \Sigma)$ and $p \in \mathbf{P}$, let $u \cdot p$ denote $\int_{\Omega} u dp$.

The following notation will be convenient for stating the first lemma. For every $a, b \in A$ let

$$Y^{ab} \equiv \{p \in \mathbf{P} \mid a \succ_p b\} \text{ and}$$

$$W^{ab} \equiv \{p \in \mathbf{P} \mid a \%_p b\}.$$

Observe that by definition and A1: $Y^{ab} \subset W^{ab}$, $W^{ab} \cap Y^{ba} = \emptyset$, and $W^{ab} \cup Y^{ba} = \mathbf{P}$. The first main step in the proof of the theorem is:

Lemma 1 *There exists $\{u^{ab}\}_{a,b \in A, a \neq b} \subset \mathbf{B}(\Omega, \Sigma)$ such that, for every distinct $a, b \in A$:*

- (i) $W^{ab} = \{p \in \mathbf{P} \mid u^{ab} \cdot p \geq 0\}$;
- (ii) $Y^{ab} = \{p \in \mathbf{P} \mid u^{ab} \cdot p > 0\}$;
- (iii) $W^{ba} = \{p \in \mathbf{P} \mid u^{ab} \cdot p \leq 0\}$;
- (iv) $Y^{ba} = \{p \in \mathbf{P} \mid u^{ab} \cdot p < 0\}$;
- (v) *Neither $u^{ab} \leq 0$ nor $u^{ab} \geq 0$;*
- (vi) $-u^{ab} = u^{ba}$.

Moreover, $\{u^{ab}\}_{a,b \in A, a \neq b}$ is unique in the following sense: for every distinct $a, b \in A$, u^{ab} and u^{ba} may be multiplied by an arbitrary positive constant $\lambda_{\{a,b\}} > 0$.

The lemma states that we can associate with every pair of distinct acts $a, b \in A$ a separating hyperplane defined by $u^{ab} \cdot p = 0$ ($p \in \mathbf{P}$), such that

$a \%_p b$ iff p is on a given side of the plane (i.e., iff $u^{ab} \cdot p \geq 0$). Observe that if there are only two acts, Lemma 1 completes the proof of sufficiency: for instance, one may set $u^a = u^{ab}$ and $u^b = 0$. More generally, we will show in the following lemmata that one can find a function u^a for every act a , such that, for every $a, b \in A$, u^{ab} is a positive multiple of $(u^a - u^b)$.

For a subset B of P let $\text{int}(B)$ denote the set of interior points of B (relative to P).

Proof of Lemma 1:

The combination axiom implies that the sets Y^{ba} and W^{ab} are convex. The continuity axiom implies that the sets Y^{ba} are open (in the relative topology). This, in turn, implies that the sets W^{ab} are closed and therefore compact in the weak* topology. This allows the use of a (weak) separating hyperplane theorem between two disjoint and convex sets, one of which is compact: the convex hull of W^{ab} and the origin (i.e., $\{\alpha p \mid \alpha \in [0, 1], p \in W^{ab}\}$) on the one hand, and $\{\alpha p \mid \alpha \in (0, 1], p \in Y^{ba}\}$ on the other. That is, we obtain a non-zero function $u^{ab} \in \mathbf{B}(\Omega, \Sigma)$ such that $u^{ab} \cdot p \geq 0$ for all $p \in W^{ab}$ and $u^{ab} \cdot p \leq 0$ for all $p \in Y^{ba}$. Further, we argue that u^{ab} does not vanish on $W^{ab} \cup Y^{ba} = \mathbf{ba}_+^1(\Omega, \Sigma) = P$. If it did, then it would also vanish on $\mathbf{ba}_+(\Omega, \Sigma)$, and therefore also on $\mathbf{ba}_-(\Omega, \Sigma)$. But in this case it would vanish on all of $\mathbf{ba}(\Omega, \Sigma)$, in view of Jordan's decomposition theorem, in contradiction to the fact that u^{ab} is non-zero.

We argue that for some $p \in Y^{ba}$, $u^{ab} \cdot p < 0$. If not, $u^{ab} \cdot p = 0$ for all $p \in Y^{ba}$. Since u^{ab} does not vanish on $W^{ab} \cup Y^{ba}$, there has to exist a $q \in W^{ab}$ with $u^{ab} \cdot q > 0$. But then for all $\varepsilon > 0$, $u^{ab} \cdot (\varepsilon q + (1 - \varepsilon)p) > 0$, while $\varepsilon q + (1 - \varepsilon)p \in Y^{ba}$ for small enough ε by the continuity axiom. Next, we argue that for all $q \in Y^{ab}$ we have $u^{ab} \cdot q > 0$. Indeed, if $u^{ab} \cdot q = 0$ for $q \in Y^{ab}$, $u^{ab} \cdot (\varepsilon p + (1 - \varepsilon)q) < 0$ for all $\varepsilon > 0$. By a similar argument, $u^{ab} \cdot p < 0$ for all $p \in Y^{ba}$.

Thus $Y^{ba} \subset \{p \mid u^{ab} \cdot p < 0\}$. Since we also have $W^{ab} \subset \{p \mid u^{ab} \cdot p \geq 0\}$, $Y^{ba} \supset \{p \mid u^{ab} \cdot p < 0\}$. That is, $Y^{ba} = \{p \mid u^{ab} \cdot p < 0\}$ and $W^{ab} = \{p \mid u^{ab} \cdot p \geq 0\}$.

0}. We have also shown that $Y^{ab} \subset \{p \mid u^{ab} \cdot p > 0\}$. To show the converse inclusion, assume that $u^{ab} \cdot p > 0$ but $a \sim_p b$. Choose $q \in Y^{ba}$. By the combination axiom, $\alpha p + (1 - \alpha)q \in Y^{ba}$ for all $\alpha \in (0, 1)$. But for α close enough to 1 we have $u^{ab} \cdot (\alpha p + (1 - \alpha)q) > 0$, a contradiction. Hence $Y^{ab} = \{p \mid u^{ab} \cdot p > 0\}$ and $W^{ba} = \{p \mid u^{ab} \cdot p \leq 0\}$.

Observe that u^{ab} can be neither non-positive nor non-negative due to the diversity axiom (applied to the pair a, b).

We now turn to prove uniqueness. Assume that $u^{ab}, v^{ab} \in \mathbf{B}(\Omega, \Sigma)$ both satisfy conditions (i)-(v) of Lemma 1. Consider a two-person zero-sum game with a payoff matrix $(u^{ab}, -v^{ab})$. Specifically, (i) the set of pure strategies of player 1 (the row player) is Ω ; (ii) the set of pure strategies of player 2 (the column player) is $\{L, R\}$; and (iii) if player 1 chooses $\omega \in \Omega$, and player 2 chooses L , the payoff to player 1 will be $u^{ab}(\omega)$, whereas if player 2 chooses R , the payoff to player 1 will be $-v^{ab}(\omega)$. Since both u^{ab}, v^{ab} satisfy conditions (i)-(iv), there is no $p \in \mathbf{P}$ for which $u^{ab} \cdot p > 0, -v^{ab} \cdot p > 0$. Hence the maximin in this game is non-positive. Therefore, so is the minimax. It follows that there exists a mixed strategy of player 2 that guarantees a non-positive payoff against any pure strategy of player 1. In other words, there are $\alpha, \beta \geq 0$ with $\alpha + \beta = 1$ such that $\alpha u^{ab}(\omega) \leq \beta v^{ab}(\omega)$ for all $\omega \in \Omega$. Moreover, by condition (v) $\alpha, \beta > 0$. Hence for $\gamma = \beta/\alpha > 0$, $u^{ab} \leq \gamma v^{ab}$. Applying the same argument to the game $(-u^{ab}, v^{ab})$, we find that there exists $\delta > 0$ such that $u^{ab} \geq \delta v^{ab}$. Therefore, $\gamma v^{ab} \geq u^{ab} \geq \delta v^{ab}$ for $\gamma, \delta > 0$. In view of part (v), there exists $\omega \in \Omega$ with $v^{ab}(\omega) > 0$, implying $\gamma \geq \delta$. By the same token there exists $\omega' \in \Omega$ with $v^{ab}(\omega') < 0$, implying $\gamma \leq \delta$. Hence $\gamma = \delta$ and $u^{ab} = \gamma v^{ab}$.

Finally, we prove part (vi). Observe that both u^{ab} and $-u^{ba}$ satisfy (i)-(iv) (stated for the ordered pair (a, b)). By the uniqueness result, $-u^{ab} = \alpha u^{ba}$ for some positive number α . At this stage we redefine the functions $\{u^{ab}\}_{a,b \in A}$ from the separation result as follows: for every unordered pair $\{a, b\} \subset A$ one of the two ordered pairs, say (b, a) , is arbitrarily chosen and then u^{ab} is

rescaled such that $u^{ab} = -u^{ba}$. (If A is of an uncountable power, the axiom of choice has to be used.) \square

Lemma 2 *For every three distinct acts, $f, g, h \in A$, and the corresponding vectors u^{fg}, u^{gh}, u^{fh} from Lemma 1, there are unique $\alpha, \beta > 0$ such that:*

$$\alpha u^{fg} + \beta u^{gh} = u^{fh} .$$

The key argument in the proof of Lemma 2 is that, if u^{fh} is not a linear combination of u^{fg} and u^{gh} , one may find $p \in \mathbb{P}$ for which \succ_p is cyclical. The detailed proof follows closely that of the corresponding lemma in Gilboa and Schmeidler (1999).

If there are only three acts $f, g, h \in A$, Lemma 2 allows us to complete the proof as follows: choose a function $u^{fh} \in \mathbf{B}(\Omega, \Sigma)$ that separates between f and h . Then choose the multiples of u^{fg} and of u^{gh} defined by the lemma, and proceed to define $u^f = u^{fh}$, $u^g = \beta u^{gh}$, and $u^h = 0$.

If there are more than three acts, Lemma 2 shows that the ranking of every triple of acts can be represented as in the theorem. The question that remains is whether these separate representations (for different triples) can be “patched” together in a consistent way.

Lemma 3 *There are functions $\{u^{ab}\}_{a,b \in A, a \neq b} \subset \mathbf{B}(\Omega, \Sigma)$, as in Lemma 1, such that, for any three distinct acts, $f, g, h \in A$, the Jacobi identity $u^{fg} + u^{gh} = u^{fh}$ holds.*

The proof is by induction, which is transfinite if A is uncountably infinite. The main idea of the proof is the following. Assume that one has rescaled the functions u^{ab} for all acts a, b in some subset of acts $X \subset A$, and one now wishes to add another act to this subset, $d \notin X$. Choose $a \in X$ and consider the functions u^{ad}, u^{bd} for $b \in X$. By Lemma 2, there are unique positive coefficients α, β such that $u^{ab} = \alpha u^{ad} + \beta u^{db}$. One would like to show that the coefficient α does not depend on the choice of $b \in X$. Indeed, if it did, one would find that there are $a, b, c \in X$ such that the vectors u^{ad}, u^{bd}, u^{cd} are

linearly dependent, and this contradicts the diversity axiom. Again, details are to be found in Gilboa and Schmeidler (1997, 1999).

Note that Lemma 3, unlike Lemma 2, guarantees the possibility to rescale *simultaneously* all the u^{ab} -s from Lemma 1 such that the Jacobi identity will hold on A .

We now complete the proof that (i) implies (ii). Choose an arbitrary act, say, e in A . Define $u^e = 0$, and for any other act, a , define $u^a = u^{ae}$, where the u^{ae} -s are from Lemma 3.

Given $p \in \mathbf{P}$ and $a, b \in A$ we have:

$$\begin{aligned} a \succ_p b &\Leftrightarrow u^{ab} \cdot p \geq 0 \Leftrightarrow (u^{ae} + u^{eb}) \cdot p \geq 0 \Leftrightarrow \\ &(u^{ae} - u^{be}) \cdot p \geq 0 \Leftrightarrow u^a \cdot p - u^b \cdot p \geq 0 \Leftrightarrow u^a \cdot p \geq u^b \cdot p \end{aligned}$$

Defining $u(a, \cdot) = u^a(\cdot)$, (**) of the theorem has been proved.

It remains to be shown that the functions defined above form a diversified matrix. This follows from the following results (adapted from Gilboa and Schmeidler (1999, revised version – 2001)):

Proposition 2 *Let Y be a set. Assume first $|A| \geq 4$. A matrix $u : X \times Y \rightarrow \mathbf{R}$ is diversified iff for every list (a, b, c, d) of distinct elements of A , the convex hull of differences of the row-vectors $(u(a, \cdot) - u(b, \cdot))$, $(u(b, \cdot) - u(c, \cdot))$, and $(u(c, \cdot) - u(d, \cdot))$ does not intersect \mathbf{R}_-^Y . Similar equivalence holds for the case $|A| < 4$.*

Lemma 4 *For every list (a, b, c, d) of distinct elements of A ,*

$$\text{conv}(\{u^{ab}, u^{bc}, u^{cd}\}) \cap \mathbf{R}_-^\Omega = \emptyset .$$

iff there exists $p \in \mathbf{P}$ such that $a \succ_p b \succ_p c \succ_p d$.

This completes the proof that (i) implies (ii). \square

Part 2: (ii) implies (i)

It is straightforward to verify that if $\{\%_p\}_{i \in P}$ are representable by $\{u(a, \cdot)\}_{a \in A} \subset \mathbf{B}(\Omega, \Sigma)$ as in (*), they have to satisfy Axioms 1-3. To show that Axiom 4 holds, we quote Lemma 4 and Proposition 2 of the previous part. \square

Part 3: Uniqueness

Similar to the proof in Gilboa and Schmeidler (1997,1999). $\square \square$

References

- Ashkenazi, G. and E. Lehrer (2001), “Well-Being Indices”, mimeo.
- Aumann, R. J. and A. Brandenburger (1995), “Epistemic Conditions for Nash Equilibrium,” *Econometrica*, 63: 1161-1180.
- Dreze, J. H. (1961), “Le fondements logique de l’utilite cardinale et de la probabilité subjective”, La Decision. Paris: CRNS. Translated and reprinted as “Decision Theory with Moral Hazard and State-Dependent Preferences”, in Dreze, J. H. (1987), *Essays on Economic Decisions under Uncertainty*. Cambridge: Cambridge University Press.
- Elster, J. (1998), “Emotions and Economic Theory”, *Journal of Economic Literature*, 36: 47-74.
- Fishburn, P. C. (1976), “Axioms for Expected Utility in n -Person Games”, *International Journal of Game Theory* 5: 137-149.
- Fishburn, P. C. and F. S. Roberts (1978), “Mixture Axioms in Linear and Multilinear Utility Theories”, *Theory and Decision* 9: 161-171.
- Frank, R. H. (1988). *Passions Within Reason: The Strategic Role of the Emotions*. WW Norton.
- Gilboa, I., and D. Schmeidler (1997), “Act Similarity in Case-Based Decision Theory”, *Economic Theory* 9: 47-61.
- Gilboa, I. and D. Schmeidler (1999), “Inductive Inference: An Axiomatic Approach”, forthcoming, *Econometrica*.
- Gilboa, I. and D. Schmeidler (2001), *A Theory of Case-Based Decisions*. Cambridge: Cambridge University Press.
- Gul, F. (1991), “A Theory of Disappointment Aversion,” *Econometrica*, 59: 667-686.
- Guth, W. and R. Tietz (1990), “Ultimatum Bargaining Behavior: A Survey and Comparison of Experimental Results”, *Journal of Economic Behavior and Organization*, 11: 417-449.
- Hammond, P.J. “Consequentialism and Bayesian Rationality in Normal Form Games”, in W. Leinfellner and E. Kohler (eds.) *Game Theory, Experi-*

ence, Rationality. Foundations of Social Sciences, Economics and Ethics. In Honor of John C. Harsanyi. (Vienna Circle Institute Yearbook 5) (Kluwer Academic Publishers, 1997).

Karni, E., D. Schmeidler, K. Vind (1983), "On State Dependent Preferences and Subjective Probabilities", *Econometrica* 51: 1021-1031.

Loewenstein, G. (2000), "Emotions in economic theory and economic behavior", *American Economic Review: Papers and Proceedings*, 90: 426-432.

Loomes, G. and R. Sugden (1982), "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty", *Economic Journal*, 92: 805-824.

Luce, R. D. and H. Raiffa (1957), "Games and Decisions". New York: Wiley.

Rabin, M. (1998), "Psychology and Economics", *Journal of Economic Literature*, 36: 11-46.

Roth, A. E. (1992), "Bargaining Experiments" in J. Kagel and A. E. Roth (eds.), *Handbook of Experimental Economics*, Princeton: Princeton University Press.

Rubinstein, A. (2000), *Economics and Language*. Cambridge: Cambridge University Press.

von Neumann, J. and O. Morgenstern (1944), *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.