JOURNAL OF ECONOMIC THEORY 55, 229–246 (1991)

Existence and Dynamic Consistency of Nash Equilibrium with Non-expected Utility Preferences*

EDDIE DEKEL

Department of Economics, University of California, Berkeley, California 94720

ZVI SAFRA

University of Pennsylvania, Philadelphia, Pennsylvania 19104, and Tel Aviv University, Tel Aviv, Israel

AND

UZI SEGAL

University of Toronto, Toronto, Ontario, Canada

Received March 6, 1990; revised December 31, 1990

Sufficient conditions for the existence of a Nash equilibrium are given when preferences may violate the reduction of compound lotteries assumption (RCLA). Without RCLA decision makers may not be indifferent between compound lotteries which have the same probabilities of final outcomes. Therefore the conditions depend on how players perceive the game—whether they view themselves as moving first or second. We also review conditions under which the equilibria will be dynamically consistent. *Journal of Economic Literature* Classification Number: 026.

I. INTRODUCTION

There have recently been many generalizations of expected utility (EU) theory.[1] These generalizations are concerned with the robustness of the EU model, and with explaining data which conflicts with the EU hypothesis. In order for such generalizations to be applied in economic models there are

* The authors thank Edi Karni, Ariel Rubinstein, David Schmeidler, and Menahem Yaari for helpful comments and the National Science Foundation and SSHRC for financial support.

[1] A seminal reference is Machina [19]. For surveys of the literature, including empirical findings and generalizations of EU, see, for example, Macrimmon and Larson [24], Machina [20–22], and Fishburn [7].

0022-0531/91 $3.00

at least two issues which must be addressed.[2] First, how should interactions among non-EU decision-makers be modelled, that is, how is game theory to be extended to non-EU preferences? Second, how are decisions over time to be modelled, in particular, when are non-EU decision-makers dynamically consistent? The former question has been addressed by Crawford [4] and Karni and Safra [17], the latter by Green [11], Gul and Lantto [13], Hammond [14], and Machina [23], among others. Regarding existence, Crawford [4] and Karni and Safra [17] show that, in general, Nash equilibrium may fail to exist (in particular in the absence of quasi-concavity of preferences over lotteries). However, Crawford shows that existence of an equilibrium in *beliefs* can be guaranteed (the reader is referred to Crawford's insightful paper for details regarding this result and the reason for a distinction between equilibrium in beliefs and in strategies, cf. Aumann [1]). The analysis of dynamic consistency has used a variety of assumptions and different methods of applying preferences in dynamic models. Depending on the assumptions invoked, results have ranged from the conclusion that dynamic consistency imposes no restrictions on preferences (Machina [23]), through the restriction that it implies quasi-convexity (Green [11]), quasi-concavity *and* quasi-convexity (i.e., betweenness) (Gul and Lantto [13]), to the conclusion that only expected utility preferences satisfy dynamic consistency (e.g., Hammond [14] and Karni and Safra [18]).

To clarify the issue of dynamic consistency, consider a decision-maker who can choose between two lotteries, X or Y. Suppose that this person is indifferent between X and Y, but prefers throwing a fair coin (denoted by $Z = 0.5\,X + 0.5\,Y$), see Machina [23, Sect. 4.2]. Hence $V(Z) > V(X) = V(Y)$, where $V(\cdot)$ is the individual's preference functional over probability distributions. Although *ex ante* the decision-maker would prefer to throw the coin and carry out the plan of X if heads, Y if tails, after the fact one might say that the decision-maker prefers to toss the coin again. The decision-maker's dilemma appears even worse if $V(Z) > V(X) > V(Y)$, so that *ex post* Y would never be chosen (unless the *ex post* and *ex ante* preferences differed). Machina [23] addresses this issue by arguing that the *ex post* preferences are derived from *ex ante* preferences by conditioning on any borne risk, so that *ex post* and *ex ante* choices will always coincide. Karni and Safra [17] examine when decision-makers would carry out their plans, assuming that the same preference function V is used both *ex ante* and *ex post*, so that restrictions are necessary to avoid situations such as the one described above.

All of the above papers (on games and on dynamic consistency) impose

[2] Machina [23] provides a comprehensive discussion of the research which needs to be accomplished in order for non-EU models to be adopted.

the reduction of compound lotteries assumption (RCLA), so that compound lotteries (inherent in games and in dynamic choice problems) are reduced to single lotteries by multiplying the probabilities. Von Neumann and Morgenstern [25, p. 632] discuss the crucial nature of this assumption and the desirability of relaxing it:

> It seems probable, that the really critical group of axioms is (3 : C), or more specifically, the axiom (3 : C : b). This axiom expresses the combination rule for multiple chance alternatives [RCLA].... Some change of the system (3 : A) – (3 : C), at any rate involving the abandonment or at least a radical modification of (3 : C : b), may perhaps lead to a mathematically complete and satisfactory calculus of utilities, which allows for the possibility of a specific utility or disutility of gambling. It is hoped that a way will be found to achieve this, but the mathematical difficulties seem to be considerable.

Segal [26–28] has proposed and analyzed a model where individuals may violate RCLA. He has also shown that this approach can explain much of the empirical evidence that conflicts with the expected utility hypothesis. Moreover, dropping the RCLA addresses an interesting robustness question: What results are sensitive to the hypothesis that individuals multiply probabilities to reduce multi-stage lotteries to single lotteries?[3]

The first concern of this paper is with the effect of dropping RCLA on game theory. We examine this effect while retaining the other assumptions which in the context of RCLA would imply EU—namely compound independence (a multi-stage version of the von Neumann and Morgenstern independence axiom, see Segal [28]),[4] continuity, and the preference order axioms (completeness, transitivity, etc.). Our analysis is restricted to two-person normal-form games. Despite the restriction to normal-form games the analysis involves compound lotteries since each player separately chooses strategies. Moreover, there is a question of dynamic consistency since, unless they can commit to mixed strategies, in equilibrium each player must want to play any pure strategy which occurs with positive probability in her equilibrium mixed strategy.

In order to describe our results it is important to recall that the way in which a decision problem is described may affect the ranking of the available strategies. In particular, when RCLA is dropped the individual need not be indifferent between compound lotteries that have the same ultimate probabilities over outcomes. Therefore, whether players *perceive* themselves as moving first or second may be important. We show that if the players perceive themselves as moving first, then without further assumptions (but still retaining the continuity, compound independence,

[3] Machina [22, pp. 147–149] discusses and provides references to work on the systematic miscalculation of probabilities by individuals.

[4] The compound independence axiom implies that the *ex post* preferences determine the *ex ante* preferences—in a sense this is the converse of Machina's approach.

and preference order axioms) existence and dynamic consistency of Nash equilibrium is guaranteed. By contrast, when they perceive themselves as moving second, we need significant restrictions in order to guarantee existence. The standard assumptions used in proofs of existence are continuity of the preferences (in order to guarantee upper hemi-continuity of the best reply correspondences), and quasi-concavity of the preferences, in order to obtain convex valued best reply correspondences. We show that convexity of the best reply sets implies that the preferences are EU. So in one case, even if players fail to multiply probabilities, existence and dynamic consistency of Nash equilibria is guaranteed; while in the other case the standard method of proving existence requires the assumption of EU preferences, and hence the assumption that players *do* multiply probabilities.

The paper discusses first the question of existence, and then the issue of dynamic consistency. The results when RCLA is dropped were just described. When RCLA is assumed to hold our results focus on dynamic consistency, since existence of an equilibrium in beliefs has been discussed by Crawford, and for existence of an equilibrium in strategies quasi-concavity of preferences is the obvious and natural sufficient condition. Regarding dynamic consistency, we examine three models in which Nash equilibria in dynamically consistent strategies can be obtained. One is based on Machina [23], and we demonstrate by example the sense in which this approach conflicts with the notion of backwards induction and subgame perfection (see also Machina [23, Sect. 6.5]). An alternative approach based on Karni and Safra [16, 17] requires either restrictions on the preferences, in particular assuming quasi-convexity (see also Green [11]), or restrictions on the strategy space. The third approach is based on Segal [28] and assumes compound independence but not RCLA.

There are several other papers which relate to this work. There has been other research on non-EU preferences in the context of game theory. Crawford [4] and Karni and Safra [17, 18] were discussed above. The transitivity axiom is weakened by Fishburn and Rosenthal [8], and the continuity axiom is weakened by Fishburn [6] and by Skala [30]. These last three papers are quite different from ours since they retain the assumption that players' preferences are linear in their own mixed strategies. Geanakoplos, Pearce, and Stacchetti [10] allow for probabilities to enter in a non-linear manner, but through the beliefs about other players' beliefs and not directly through the strategies as we do. They also discuss the questions of existence (in the normal form) and backwards induction (sequential rationality) in the extensive form. Finally there are models of games where players make "mistakes" in calculating conditional probabilities (due to incorrect information processing), such as Geanakoplos [9] and Brandenburger, Dekel, and Geanakoplos [2].

II. The Model and Existence

Consider a two-person game, where players 1 (female) and 2 (male) have finite pure strategy sets S^i, and mixed strategy spaces Σ^i (throughout we use the index i for $i = 1, 2$ and j for $j \neq i$). The mixed strategy which assigns probability one to a pure strategy s^i in S^i will also be denoted by s^i. Each player has an outcome function $H^i\colon S^1 \times S^2 \to R$ and a preference functional $V^i\colon L \to R$, where L is the space of simple (finite outcome) lotteries over the reals. The functions V^i are assumed to be continuous and strictly monotone with respect to first-order stochastic dominance. (Of course the presumption of such a preference functional implies that the underlying (weak) preference relation on lotteries is transitive and complete.) The basic features of a game are thus given by $\Gamma = \{S^i, H^i, V^i\}$. This specification does not describe the functions, denoted $h^i\colon \Sigma^1 \times \Sigma^2 \to L$, by which the players reduce pairs of mixed strategies into lotteries. The standard reduction mechanism (called the reduction of compound lotteries assumption —RCLA) simply involves multiplying out the probabilities, so that the lottery induced by any pair of mixed strategies (i.e., $h^i(\sigma^1, \sigma^2)$) is the lottery which assigns probability $\sigma^1(s^1) \times \sigma^2(s^2)$ to the outcome $H^i(s^1, s^2)$. We write $H^i(\sigma^1, \sigma^2)$ for this lottery; that is, we extend H^i to a function from $\Sigma^1 \times \Sigma^2$ into L which simply multiplies out the probabilities. To save on notation we write $V^i(\sigma^1, \sigma^2)$ for $V^i(h^i(\sigma^1, \sigma^2))$, which is the function that, for any given mechanism h^i for reducing pairs of mixed strategies into lotteries, represents player i's preferences over pairs of mixed strategies. The best reply correspondences $BR^i\colon \Sigma^j \to \Sigma^i$ are defined as usual by $BR^i(\sigma^j) = \{\sigma^i \mid V^i(\sigma^i, \sigma^j) \geqslant V^i(\tau^i, \sigma^j) \ \ \forall \tau^i \in \Sigma^i\}$. Finally, a pair of mixed strategies (σ^1, σ^2) is a Nash equilibrium if $\sigma^i \in BR^i(\sigma^j)$. Clearly, sufficient conditions for the existence of a Nash equilibrium are that $V^i(\sigma^i, \sigma^j)$ is continuous (in the product topology) so that BR^i is upper hemi-continuous and non-empty valued, and that $V^i(\sigma^i, \sigma^j)$ is quasi-concave in σ^i, so that BR^i is convex valued.

In order to discuss alternative methods of reducing the pair of mixed strategies into a lottery, we must consider how the players may perceive the game. Focussing on player 1, she may perceive herself as either moving first or second. These correspond to Figs. 1 and 2, which should be understood as follows. In the former case if player 1 chooses s^1 in S^1 and player 2 chooses σ^2, then player 1 faces the lottery that gives her $H^1(s^1, s^2)$ with probability $\sigma^2(s^2)$. Using the extension of H^1 discussed above we denote this lottery by $H^1(s^1, \sigma^2)$. Hence, a choice of σ^1 gives player 1 the compound lottery which yields the lottery $H^1(s^1, \sigma^2)$ with probability $\sigma^1(s^1)$. Figure 2 corresponds to the case where player 1 perceives herself as moving second. Here, if player 1 chooses σ^1 and player 2 choose s^2, then she faces the lottery $H^1(\sigma^1, s^2)$ (which yields $H^1(s^1, s^2)$ with probability $\sigma^1(s^1)$).

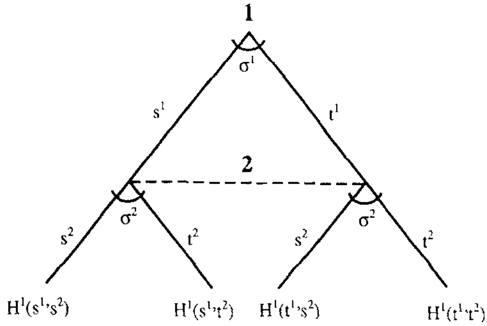

FIGURE 1

Hence, when player 2 chooses σ^2, player 1 faces the compound lottery which with probability $\sigma^2(s^2)$ yields the lottery $H^1(\sigma^1, s^2)$. We call these the first and second hypotheses on the players' perceptions of the mixed strategies, or perceptual hypotheses for short. We will also discuss two methods of evaluating compound lotteries—RCLA and compound independence (CI). Clearly under RCLA both perceptual hypotheses are equivalent to the simple lottery $H^1(\sigma^1, \sigma^2)$, hence there are three cases to consider. The example in Fig. 3 (similar to Crawford [4, Sect. 3]) demonstrates the existence problem in the case where RCLA is assumed. If V^i is strictly quasi-convex (i.e., for any X, Y with $V^i(X) \neq V^i(Y)$ and any $0 < \alpha < 1$, $V^i(\alpha X + (1-\alpha) Y)$ is strictly less than either $V^i(X)$ or $V^i(Y)$, so that, roughly speaking, players prefer not to randomize) then there is no equilibrium in mixed strategies since for any mixed strategy of the column player, the row player will strictly prefer at least one of her pure strategies. Since there is no pure strategy equilibrium, this shows the failure of existence.

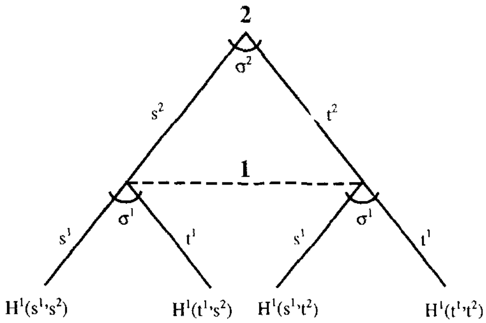

FIGURE 2

	L	R
U	1,0	0,4
D	0,1	4,0

FIGURE 3

Turning now to alternative mechanisms for reducing pairs of mixed strategies to lotteries, we introduce the axiom of compound independence which in conjunction with RCLA would imply expected utility. Compound independence is an axiom on compound lotteries, so a notation for such lotteries will prove helpful. A compound lottery is an $2n$-tuple $(X_i, p_i)_{i=1}^n$, where $X_i \in L$ and p_i is the probability of receiving lottery X_i. Let V_c represent preferences over compound lotteries, and let V represent preferences over L.

DEFINITION 1. Let $(X_i, p_i)_{i=1}^n$ and $(Y_i, p_i)_{i=1}^n$ be any pair of compound lotteries which for some $j \in \{1, ..., n\}$ satisfy $X_i = Y_i$ for all $i \neq j$. Preferences represented by V and V_c satisfy compound independence (CI) if $V_c((X_i, p_i)_{i=1}^n) \geq V_c((Y_i, p_i)_{i=1}^n)$ if and only if $V(X_j) \geq V(Y_j)$.

Given any continuous and monotonic preference functional $V: L \to R$ there is a unique continuous and monotonic function V_c from compound lotteries into R such that V and V_c satisfy CI. This function V_c is calculated as follows. Let the certainty equivalent of a lottery X be denoted by $CE(X) \in R$, so $V(X) = V(CE(X))$.[5] Given a compound lottery (X_i, p_i), let $V_c((X_i, p_i)) = V((CE(X_i), p_i))$. Henceforth, for notational simplicity, we often use V (rather than V_c) to denote the extension of V to a preference functional on compound lotteries.

The table in Fig. 4 summarizes the different methods of reducing a pair of mixed strategies into a simple lottery. We now turn to our first result on the existence of Nash equilibrium.

[5] Elements of R are used also to denote lotteries which assign probability one to that element.

	First Interpretation	Second Interpretation
RCLA	$H^1(\sigma^1,\sigma^2)$	
CI	$[CE^1(H^1(s^1,\sigma^2)),\sigma^1(s^1)]_{s^1 \in S^1}$	$[CE^1(H^1(\sigma^1,s^2)),\sigma^2(s^2)]_{s^2 \in S^2}$

FIGURE 4

THEOREM 1. *If both players satisfy the first perceptual hypothesis on mixed strategies and satisfy compound independence, then in any game Γ there exists a Nash equilibrium.*

Proof. It is sufficient to show that the best reply correspondences BR^i are upper hemi-continuous, non-empty, and convex valued. Upper hemi-continuity follows from continuity of both the preference functional V on lotteries, and of the extension V_c on compound lotteries. That the correspondences are non-empty valued follows from this continuity and the compactness of Σ^i. The main point of this proof is that *under the first perceptual hypothesis* the compound independence axiom guarantees convexity of the best reply sets. In fact $BR^i(\sigma^j)$ is equal to the convex hull of $\{s^i \in S^i\colon\ V^i(s^i,\sigma^j) \geqslant V^i(r^i,\sigma^j)$ for all $r^i \in S^i\}$. To see that $BR^i(\sigma^j)$ is a superset of this convex hull note that if any two pure strategies s^i and t^i are best replies to σ^j then $V^i(H^i(s^i,\sigma^j)) = V^i(H^i(t^i,\sigma^j))$, so the certainty equivalent of these two lotteries are equal. Therefore, the compound lottery generated by i playing any probability mixture on these two pure strategies yields the same utility as playing the pure strategies. To see that the convex hull described above is a superset of $BR^i(\sigma^j)$ it suffices to observe that if a mixed strategy σ^i is a best reply then every pure strategy in S^i assigned positive probability by σ^i is a best reply. This can be seen as follows. If $t^i \in S^i$ is not a best reply then there is an $r^i \in S^i$ such that $CE^i(H^i(r^i,\sigma^j)) \geqslant CE^i(H^i(s^i,\sigma^j))$ for all s^i, with a strict inequality for $s^i = t^i$. If t^i is assigned positive probability by σ^i then the lottery which for each s^i in S^i gives $CE^i(H^i(s^i,\sigma^j))$ with probability $\sigma^i(s^i)$ is stochastically dominated by $CE^i(H^i(r^i,\sigma^j))$. So strict monotonicity and compound independence imply that σ^i is not a best reply. Q.E.D.

Thus, under the first perceptual hypothesis, the assumption of compound independence, even without RCLA, is sufficient to guarantee existence. Moreover, it is intuitively clear, and formally stated in the next section, that under the first interpretation CI also guarantees dynamic consistency. However, the situation is quite different with the second interpretation. The example below demonstrates that CI is not sufficient to guarantee the

	L	R
U	-1,1	2,0
D	1,0	0,1

FIGURE 5

existence of a Nash equilibrium under the second interpretation. Since we assume continuity of V^i the only reason for existence to fail is that the best reply sets are not convex. This suggests asking what conditions are necessary and sufficient for convex best reply sets. The answer is that quite strong conditions are needed—under the second interpretation of the game, assuming compound independence and convex best reply sets is equivalent to assuming EU preferences.

EXAMPLE. Consider a game where the payoffs from pure strategies are as described in Fig. 5. Let player 1's preferences over simple lotteries be given by Yaari's [31] function, namely $V^1(X)=\int t\, dg(X(t))$, where X is a cumulative distribution function and

$$g(p)=\begin{cases}2p, & 0\leqslant p\leqslant 0.5\\ 1, & 0.5\leqslant p\leqslant 1.\end{cases}$$

Assume that player 2's preferences over simple lotteries are linear (i.e., V^2 is an expected utility function) so that in conjunction with compound independence his preferences over compound lotteries are also EU.[6] The best reply correspondences are as follows, where a pair $(1-p, p)$ denotes for player 1 the mixed strategy that assigns D probability p, and $(1-q, q)$ denotes for player 2 the mixed strategy that assigns R probability q.[7]

[6] Although player 1's preferences are not strictly monotone (in first order stochastic dominance) since g is not strictly increasing, this is only done for simplicity.

[7] BR^1 requires tedious computations, which are omitted.

$$BR^1((1-q,q))=\begin{cases}(0,1), & 0\leqslant q<1/2\\ \{(1-p,p)\colon p\in[3/4,1]\}, & 1/2\leqslant q<2/3\\ \{(1,0)\}\cup\{(1-p,p)\colon p\in[3/4,1]\}, & q=2/3\\ (1,0), & 2/3<q\leqslant 1\end{cases}$$

$$BR^2((1-p,p))=\begin{cases}(1,0), & 0\leqslant p<1/2\\ \text{anything}, & p=1/2\\ (0,1), & 1/2<p\leqslant 1.\end{cases}$$

These best reply correspondences are illustrated in Fig. 6, from which it is clear that the game has no Nash equilibrium.

THEOREM 2. *Assume that player i satisfies the second perceptual hypothesis and satisfies compound independence. The set $BR^i(\sigma^j)$ is convex in any game and for any opponent's strategy σ^j if and only if V^i is linear, so that i has expected utility preferences.*

Proof. "If" is trivial. For "only if," it is enough to show that for all X, Y, Z in L we have $V^i(X)=V^i(Y)\Rightarrow V^i(0.5\,X+0.5\,Z)=V^i(0.5\,Y+0.5\,Z)$ (see Herstein and Milnor [15]). Using continuity we can restrict attention to lotteries such that $X=(x_n,1/N)_{n=1}^N$, $Y=(y_n,1/N)_{n=1}^N$, $Z=(z_n,1/N)_{n=1}^N$. Consider now the game in which $H^i(s^1,s^2)$ is given by the matrix in Fig. 7. Assume that player j plays the mixed strategy σ^j which assigns each s^j probability $1/(2N)$.

First it is shown that $BR^i(\sigma^j)=\Sigma^i$, so that player i is indifferent among all her strategies. Let $\sigma^i\in BR^i(\sigma^j)$ and denote $\sigma^i(s_k^i)=\alpha_k$. We can identify any mixed strategy with probability vectors $(\beta_k)_{k=1}^{2N}$. Denote by Π, with generic element π, the set of the N permutations on $\{1,...,N\}$ of the form $(1,2,...,N)$, $(N,1\,2,...,N-1)$, $(N-1,N,1,2,...,N-2),...,(2,3,...,N,1)$. Denote by $\pi(k)$ the kth element of the permutation π. It is easy to see that for any π player i is indifferent between the mixed strategies

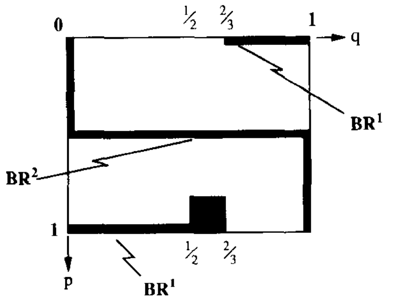

FIGURE 6

j

i	s^j_1	s^j_2	$\cdots$	s^j_N	s^j_{N+1}	s^j_{N+2}	$\cdots$	s^j_{2N}
s^i_1	x_1	x_2	$\cdots$	x_N	z_1	z_2	$\cdots$	z_N
s^i_2	x_2	x_3	$\cdots$	x_1	z_2	z_3	$\cdots$	z_1
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
s^i_N	x_N	x_1	$\cdots$	x_{N-1}	z_N	z_1	$\cdots$	z_{N-1}
s^i_{N+1}	z_1	z_2	$\cdots$	z_N	x_1	x_2	$\cdots$	x_N
s^i_{N+2}	z_2	z_3	$\cdots$	z_1	x_2	x_3	$\cdots$	x_1
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
s^i_{2N}	z_N	z_1	$\cdots$	z_{N-1}	x_N	x_1	$\cdots$	x_{N-1}

FIGURE 7

$\sigma^i_\pi \equiv (\alpha_{\pi(1)}, ..., \alpha_{\pi(N)}, \alpha_{\pi(1)+N}, ..., \alpha_{\pi(N)+N})$ and $\sigma^i \equiv (\alpha_1, ..., \alpha_{2N})$. (This is because $CE^i(H^i(\sigma^i, s^j_k)) = CE^i(H^i(\sigma^i_\pi, s^j_{\pi(k)}))$ and $CE^i(H^i(\sigma^i, s^j_{k+N})) = CE^i(H^i(\sigma^i_\pi, s^j_{\pi(k)+N}))$, $k = 1, ..., N$, so that the certainty equivalents are simply being permuted.) Similarly, player i is indifferent between $\sigma^i_{\pi N} \equiv (\alpha_{N+\pi(1)}, ..., \alpha_{N+\pi(N)}, \alpha_{\pi(1)}, ..., \alpha_{\pi(N)})$ and σ^i. Hence, for any permutation π in Π, both σ^i_π and $\sigma^i_{\pi N}$ are elements of $BR^i(\sigma^j)$. Clearly $[1/(2N), ..., 1/(2N)] = [\Sigma_\pi(\sigma^i_\pi + \sigma^i_{\pi N})]/(2N)$. By assumption $BR^i(\sigma^j)$ is convex, so that $[1/(2N), ..., 1/(2N)] \in BR^i(\sigma^j)$. Now $V^i([1/(2N), ..., 1/(2N)], \sigma^j) = V^i(0.5\, X + 0.5\, Z)$, since (under the second perceptual hypothesis) each pure strategy of player j yields the lottery $0.5\, X + 0.5\, Z$. Moreover every pure strategy of player i yields the same lottery, so (again using convexity) $BR^i(\sigma^j) = \Sigma^i$. Thus, it has been shown that player i is indifferent among all her strategies. We have also seen that there is a strategy which yields the lottery $0.5\, X + 0.5\, Z$. Now note that the mixed strategy $(1/N, ..., 1/N, 0, ..., 0)$ yields player i the compound lottery which with probability $N/(2N)$ gives X and with probability $N/(2N)$ gives Z. Therefore player i is indifferent between the lottery $0.5\, X + 0.5\, Z$ and the compound lottery which gives X and Z with probability 1/2 each. If in the above game the x_k's are replaced by y_k's then the same conclusion holds with X replaced by Y. So, $V^i(0.5\,X+0.5\,Z)=V^i_c(X, 0.5; Z, 0.5)$ and $V^i(0.5\,Y+0.5\,Z)=V^i_c(Y, 0.5; Z, 0.5)$. By hypothesis $CE^i(X)=CE^i(Y)$, so by compound independence player i is indifferent between the compound lotteries $(X, 0.5; Z, 0.5)$ and $(Y, 0.5; Z, 0.5)$. Hence, $V^i(0.5\,X+0.5\,Z)=V^i(0.5\,Y+0.5\,Z)$. Q.E.D.

Remark. It should be clear that Theorem 2 characterizes when player i's best reply correspondence is convex valued regardless of the other player's preferences. In particular if player 1 perceives herself as moving first, player 2 perceives himself as moving second, and both players satisfy compound independence, then both players best reply sets are convex for all games and opponents' strategies if and only if player 2 is an expected utility maximizer. (Convexity for player 1 follows from the proof of Theorem 1.)

Theorem 2 is somewhat weaker than what we would want since it does not show that under the second perceptual hypothesis existence is equivalent to expected utility. We do not know whether or not the stronger result can be proven. However, Theorems 1 and 2 above do yield some additional results regarding the relationship between the sets of Nash equilibria under the two approaches.

COROLLARY 1. *Assuming CI, the best reply correspondences under the first and second perceptual hypothesis are equal if and only if V^i is linear, so that player i has expected utility preferences.*

Proof. The proof of Theorem 1 implies that under the first interpretation the best reply sets are convex valued. Theorem 2 states that convex valued best reply sets imply expected utility preferences. Q.E.D.

COROLLARY 2. *Assuming CI, the set of Nash equilibria coincide under the two perceptual hypotheses if and only if the players have expected utility preferences.*

Proof. If player 1 does not have expected utility preferences, construct a game in which player 2's payoffs are constant so that he is indifferent between all his strategies. By Corollary 1 there exist payoffs for player 1 and a mixed strategy for player 2 which yield different best reply sets for player 1. Hence in the constructed game there are different Nash equilibria according to the two perceptual hypotheses. Q.E.D.

These two corollaries imply that under compound independence, if we want conclusions to be independent of the way the players perceive the game then we must assume expected utility preferences.

Remark. The extension of the analysis provided in this section to n-person games is straightforward. For example, the analog to Theorem 1 would say that if each player i would view herself as selecting among simple lotteries determined by the opponents' mixed strategies then (assuming CI) a Nash equilibrium exists. (The other players' mixed strategies can be reduced to a simple lottery in several different ways: either RCLA or CI can be used, and in the latter case player i can perceive different orders of play for her opponents. These reductions may yield different Nash equilibria, but in all cases an equilibrium exists.) Similarly, under the second perceptual hypothesis each player i would view the choice of σ^i as determining the compound lottery $[CE^i(H^i(\sigma^i, s^{-i})), \sigma^{-i}(s^{-i})]$, where σ^{-i} is a probability distribution on the opponents pure strategy combinations and s^{-i} is an index for these pure strategy combinations. In this case a result analogous to Theorem 2 would hold: Assuming CI and convex best reply sets is equivalent to assuming expected utility preferences.

III. Dynamic Consistency

In this section we examine the issue of dynamic consistency. To formalize this notion we must separate the individual's choice problem into two stages—first the choice of mixed strategy, and then the decision of whether to play the pure strategy which the mixed strategy specifies after the randomizing device is used. The main conclusion of this section is that there are several interesting ways of modelling the game which yield the existence of Nash equilibria that are dynamically consistent.

Let $W^i: \Sigma^j \times \Sigma^i \times S^i \times S^i \to R$ be the individual's preference functional at the interim stage. That is, it represents preferences for a given mixed strategy of the opponent (the first argument), after having chosen a mixed strategy (the second argument), and after having observed the outcome of one's own randomization (the third argument), over one's choice of a pure strategy (the fourth argument). The relationship between the third and fourth arguments is the same as that between a mediator's recommendation and the actual choice in a correlated equilibrium. Dynamic consistency is the requirement that the player actually play according to the mixed strategy (the mediator's recommendation).

Definition 2. Given a mixed strategy of the opponent σ^j, a mixed strategy σ^i is dnamically consistent if for all s^i such that $\sigma^i(s^i) > 0$, $W^i(\sigma^j, \sigma^i, s^i, s^i) \geqslant W^i(\sigma^j, \sigma^i, s^i, t^i)$ for all t^i in S^i.

Thus, *ex ante* the player chooses σ^i according to V^i, and at the "interim" stage—having observed the outcome of her randomizing device—she chooses $s^i \in S^i$ according to W^i. The substantive issue concerns the relationship between W^i and V^i.[8] To analyze this issue we assume that we are given a basic preference over lotteries $U^i: L \to R$, and describe how V^i and W^i can be determined from U^i.

Model 1. This approach is based on Machina's [23] idea that preferences are affected by risks born. In the context considered here this can be formalized as follows. Let $V_1^i(\sigma^i, \sigma^j) = U^i(H^i(\sigma^i, \sigma^j))$; and for r, t in S^i let $W_1^i(\sigma^j, \sigma^i, r, t) = V_1^i(\sigma^i_{r+t}, \sigma^j)$, where

$$\sigma^i_{r+t}(s^i) = \begin{cases} \sigma^i(s^i), & s^i \neq r, t \\ \sigma^i(t) + \sigma^i(r), & s^i = t \\ 0, & s^i = r. \end{cases}$$

The interpretation of W_1^i is the following. Having decided upon σ^i *ex ante*, the interim utility of playing t instead of r (in the event that r should be played according to σ^i) is the same as having originally decided upon playing t instead of r. Therefore if at the interim stage a player would have preferred not to follow the *ex ante* choice, then the *ex ante* choice could not have been optimal. This leads to the following proposition.

PROPOSITION 1. *Let V_1^i, W_1^i be determined from U^i as in Model* 1 *above. If (σ^1, σ^2) is a Nash equilibrium then σ^i is dynamically consistent given σ^j.*

Proof. $V_1^i(\sigma^i, \sigma^j) \geqslant V_1^i(\tau^i, \sigma^j)$ for all $\tau^i \in \Sigma^i$ implies,, by definition, $W_1^i(\sigma^j, \sigma^i, s^i, s^i) \geqslant W_1^i(\sigma^j, \sigma^i, s^i, t^i)$ for all s^i such that $\sigma^i(s^i) > 0$. Q.E.D.

Despite this result, it is worth emphasizing that this approach conflicts with the idea of backwards induction (see also Machina [23]). Consider the extensive-form game in Fig. 8, where player 1's payoffs are in utiles and she has EU preferences, and player 2's preferences over mixtures over the lotteries X, Y, Z are described in the simplex of Fig. 9. This is essentially a game of perfect information. Nevertheless there are three trembling hand perfect equilibria (Selten [29]): (R, r); (L, l); and $[(L, \alpha; R, 1-\alpha), (l, 1/2; r, 1/2)]$. Moreover, since Machina's approach requires that preferences over future choices are determined from the *ex ante* preferences it seems natural to say that these are subgame perfect. (Alternatively stated, Machina's approach and definition of interim preferences does not permit "snipping the decision tree," while the intuition and definition of subgame

[8] It is important to emphasize that V^i represents the players' preferences over mixed strategies, and W^i represents preferences over final choices (which are assumed to be elements of S^i). One could imagine adding intermediate stages where the player would re-evaluate choices. However, it seems that there is only one point in time at which such a re-evaluation is meaningful—when the choice must be made.

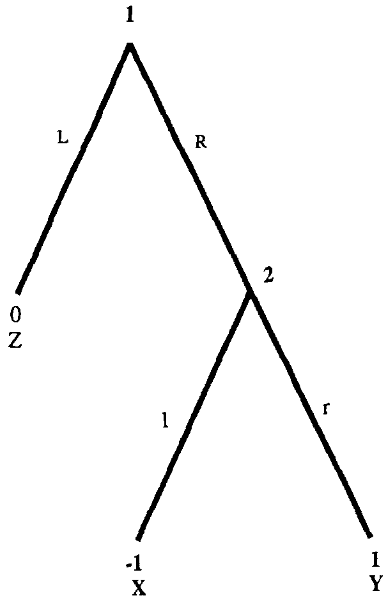

FIGURE 8

perfection *is* based on considering subgames separately. Therefore, the idea of subgame perfection conflicts with Machina's approach.)[9] In a recent paper Faruk Gul [12] made similar points using a related example. His paper focusses on various difficulties inherent in non-EU theories, with an emphasis on dynamic contexts. This example also helps relate this paper to Geanakoplos, Pearce, and Stacchetti [10] who have shown that non-linear preferences can yield multiple subgame perfect equilibria (although in their paper trembling hand perfect equilibria may fail to exist).

Model 2 (Karni and Safra [17]). $V_2^i(\sigma^i, \sigma^j) = U^i(H^i(\sigma^i, \sigma^j))$; $W_2^i(\sigma^j, \sigma^i, s^i, t^i) = V_2^i(t^i, \sigma^j)$. In this model the preferences at both stages are given by U^i.

PROPOSITION 2. *Let V_2^i, W_2^i be determined as above. If (σ^1, σ^2) is a Nash equilibrium and if U^i is quasi-convex, then σ^i is dynamically consistent given σ^j.*

Proof. If U^i is quasi-convex and $V_2^i(\sigma^i, \sigma^j) \geqslant V_2^i(\tau^i, \sigma^j)$ for all τ^i in Σ^i then for every s^i such that $\sigma^i(s^i) > 0$, $V_2^i(s^i, \sigma^j) \geqslant V_2^i(\tau^i, \sigma^j)$ for each $\tau^i \in \Sigma^i$.

[9] A slight modification of this example shows that trembling hand perfection may be more selective: change player 2's preferences in the neighborhood of Z so that indifference curves are upward sloping. Then (L, l) is no longer trembling hand perfect.

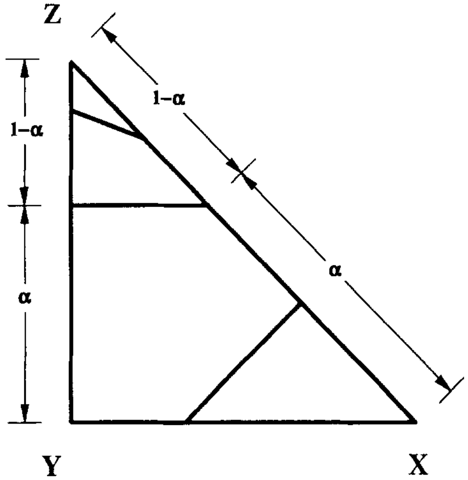

FIGURE 9

Hence, for every s^i such that $\sigma^i(s^i) > 0$, $W_2^i(\sigma^j, \sigma^i, s^i, s^i) \geqslant W_2^i(\sigma^j, \sigma^i, s^i, t^i)$ for every $t^i \in S^i$. Q.E.D.

In general, however, a Nash equilibrium may fail to be dynamically consistent. To see this consider an equilibrium (σ^1, σ^2) with strictly quasi-concave preferences where $V_2^i(s^i, \sigma^j)$ is not equal for all s^i in the support of σ^i. Clearly in such a case the equilibrium is not dynamically consistent. To address a related problem Karni and Safra [17] proposed that players will restrict attention to dynamically consistent strategies. That means that players, when deciding which strategy σ^i to choose *ex ante* in response to a mixed strategy σ^j of their opponent, will only consider strategies which are dynamically consistent given σ^j. Let $DC^i(\sigma^j)$ be the set of strategies of i that are dynamically consistent given σ^j. Then $DC^i(\sigma^j)$ and $BR^i(\sigma^j)$ are

$$DC^i(\sigma^j) = \{\sigma^i : \sigma^i(s^i) > 0 \Rightarrow V_2^i(s^i, \sigma^j) \geqslant V_2^i(t^i, \sigma^j) \text{ for all } t^i \text{ in } S^i\}.$$

$$BR^i(\sigma^j) = \{\sigma^i \in DC^i(\sigma^j) : V_2^i(\sigma^i, \sigma^j) \geqslant V_2^i(\tau^i, \sigma^j) \text{ for all } \tau^i \in DC^i(\sigma^j)\}.$$

If U^i (and hence V_2^i) is quasi-concave the best reply correspondence is clearly convex valued. However, upper-hemi-continuity of BR^i may fail in the absence of quasi-convexity. Therefore, even if the strategies are restricted to be dynamically consistent, one is led to assuming quasi-convexity and quasi-concavity (i.e., betweenness, see Chew [3], Dekel [5]) in order to guarantee existence of a dynamically consistent Nash equilibrium.

Model 3 (Segal [28]). The approach using CI is the same as discussed in the previous section. We restrict attention to the first perceptual hypothesis.

In this case $W_3^i(\sigma^j, \sigma^i, t^i, s^i) = U^i(H^i(s^i, \sigma^j))$. To define V_3^i let $h^i(\sigma^i, \sigma^j)$ be the lottery that assigns probability $\sigma^i(s^i)$ to $CE^i(H^i(s^i, \sigma^j)) \in R$ and set $V_3^i(\sigma^i, \sigma^j) = U^i(h(\sigma^i, \sigma^j))$.

PROPOSITION 3. *Let V_3^i, W_3^i be determined as above. If (σ^1, σ^2) is a Nash equilibrium then σ^i is dynamically consistent given σ^j.*

Proof. Follows from the proof of Theorem 1.

We have presented three models in which a Nash equilibrium in dynamically consistent strategies can be obtained. In the model based on Machina [23] the first period preferences are given by U, and second period preferences are induced from the first period preferences in a way that guarantees dynamic consistency. The model based on Segal [28] does the opposite: second period preferences are given by U and the first period preferences are induced from the second period preferences in a way that guarantees dynamic consistency (i.e., using compound independence). In the model based on Karni and Safra [17] both the first and the second period preferences were given by U. In this case restrictions (either on preferences or on the strategy space) are necessary to guarantee dynamic consistency.

REFERENCES

1. R. AUMANN, Correlated equilibrium as an expression of Bayesian rationality, *Econometrica* **55** (1987), 1–18.
2. A. BRANDENBURGER, E. DEKEL, AND J. GEANAKOPLOS, "Correlated Equilibrium with Generalized Information Structures," WP No. 89-010, Harvard Business School, 1989.
3. S. H. CHEW, Axiomatic utility theories with betweenness property, *Ann. Oper. Res.* **19** (1989), 273–298.
4. V. CRAWFORD, Equilibrium without independence, *J. Econ. Theory* **50** (1990), 127–154.
5. E. DEKEL, An axiomatic characterization of preferences under uncertainty: Weakening the independence axiom, *J. Econ. Theory* **40** (1986), 304–318.
6. P. C. FISHBURN, On the foundations of game theory: The case of non-archimedean utilities, *Int. J. Game Theory* **1** (1972), 65–71.
7. P. C. FISHBURN, "Nonlinear Preference and Utility Theory," Johns Hopkins Univ. Press, Baltimore, 1988.
8. P. C. FISHBURN AND R. ROSENTHAL, Non-cooperative games and nontransitive preferences, *Math. Social Sci.* **12** (1986), 1–7.
9. J. GEANAKOPLOS, "Game Theory without Partitions," DP No. 914, Cowles Foundation, Yale University, 1989.
10. J. GEANAKOPLOS, D. PEARCE, AND E. STACCHETTI, Psychological games and sequential rationality, *Games Econ. Behav.* **1** (1989), 60–79.
11. J. R. GREEN, "Making book against oneself," the independence axiom and non-linear utility theory, *Quart. J. Econ.* **102**, No. 4 (1987), 785–796.
12. F. GUL, Why do economists resist non-expected utility preferences? mimeo, Graduate School of Business, Stanford University, 1990.

13. F. GUL AND O. LANTTO, Betweenness satisfying preferences and dynamic choice, mimeo, Graduate School of Business, Stanford University, 1989.
14. P. HAMMOND, Consequentialist foundations for expected utility, *Theory Decision* **25**, No. 1 (1988), 25–78.
15. I. HERSTEIN AND S. MILNOR, An axiomatic approach to measurable utility, *Econometrica* **21** (1953), 291–297.
16. E. KARNI AND Z. SAFRA, Behaviorally consistent optimal stopping rules, *J. Econ. Theory* **51** (1990), 391–402.
17. E. KARNI AND Z. SAFRA, Ascending bid auctions with behaviorally consistent bidders, *Ann. Oper. Res.* **19** (1989), 435–446.
18. E. KARNI AND Z. SAFRA, Dynamic consistency, revelation in auctions, and the structure of preferences, *Rev. Econ. Stud.* **56** (1989), 421–434.
19. M. J. MACHINA, Expected utility analysis without the independence axiom, *Econometrica* **50** (1982), 277–323.
20. M. J. MACHINA, Generalized expected utility analysis and the nature of observed violations of the independence axiom, *in* "Foundations of Utility and Risk Theory with Application" (Bernt Stigur and Fred Wenstp, Eds.), Reidel, Dordrecht, 1983.
21. M. J. MACHINA, "The Economic Theory and Individual Behavior and Risk: Theory, Evidence and New Directions," IMSSS Tech. Rep., Stanford University, 1983.
22. M. J. MACHINA, Choice under uncertainty: Problems solved and unsolved, *J. Econ. Perspectives* **1** (1987), 121–154.
23. M. J. MACHINA, Dynamic consistency and non-expected utility models of choice under uncertainty, *J. Econ. Lit.* **27**, No. 4 (1989), 1622–1668.
24. H. MACRIMMON AND S. LARSON, Utility theory: Axioms versus paradoxes, *in* "Expected Utility Hypothesis and the Allais Paradox" (M. Allais and O. Hagen, Eds.), Reidel, Dordrecht, 1979.
25. J. VON NEUMANN AND O. MORGENSTERN, "Theory of Games and Economic Behavior," 3rd ed., Wiley, New York, 1953.
26. U. SEGAL, The Ellsberg paradox and risk aversion: An anticipated utility approach, *Int. Econ. Rev.* **28** (1987), 175–202.
27. U. SEGAL, Does the preference reversal phenomenon necessarily contradict the independence axiom? *Amer. Econ. Rev.* **78** (1988), 233–236.
28. U. SEGAL, Two stage lotteries without the reduction axiom, *Econometrica* **58** (1990), 349–377.
29. R. SELTEN, Reexamination of the perfectness concept of equilibrium points in extensive games, *Int. J. Game Theory* **4** (1975), 25–55.
30. H. J. SKALA, Nonstandard utilities and the foundations of game theory, *Int. J. Game Theory* **3** (1989), 67–81.
31. M. YAARI, The dual theory of choice under risk, *Econometrica* **55**, No. 1 (1987), 95–115.