

Discussion Paper No. 988

CREDIBLE EQUILIBRIA IN GAMES WITH
UTILITIES CHANGING DURING THE PLAY*

by

J. L. Ferreira**

I. Gilboa***

and

M. Maschler****

February 1992

*We wish to express our thanks and gratitude to Chaim Fershtman and Ehud Kalai for several conversations, from which we benefitted greatly. In particular, Chaim Fershtman pointed out to us that games with changing utilities require further analysis, thus originating this research. Ehud Kalai, among other good suggestions, insisted that we verify Theorem 4.1. This research was done while the first author was a student and the third was a visitor at KGSM-MEDS, Northwestern University. The authors express their gratitude to Northwestern University for its hospitality. The second author gratefully acknowledges NSF grant no. SES-9113108.

** University of Pennsylvania, Philadelphia, PA

*** Northwestern University, Evanston, IL

**** The Hebrew University of Jerusalem

ABSTRACT

Whenever one deals with an interactive decision situation of long duration, one has to take into account that priorities of the participants may change during the conflict. In this paper we propose an extensive-form game model to handle such situations and suggest and study a solution concept, called *credible equilibrium*, which generalizes the concept of Nash equilibrium. We also discuss possible variants to this concept and applications of the model to other types of games.

1. Introduction

As far as we know, a game in which priorities change during the play was first discussed by Homer in the story of Odysseus/Ulysses and the sirens [Odyssey, c 10th century B.C.]. According to the story, the singing of the sirens was renowned for its beauty, and so seductive that whoever heard it was lured to his death.

Odysseus achieved his desire to hear the singing of the sirens in a way that game theorists, often unjustifiably, ignore. He followed the principle¹ that when people do not like the rules of the game, they change the rules. Thus Odysseus invented a new, pure strategy; namely, he put wax in the ears of each of his sailors, commanding them to tie him to the mast and not to release him until they were safely bound for home.

Lacking the creativity and imagination of Homer, we shall not follow this avenue and regard our games as given and fixed.

Classical economics and game theory assume that an individual's preferences are constant throughout the decision-making process, even if the latter has several stages. This seemingly implausible assumption is of great theoretical value; in particular, it allows a multi-stage decision problem to be "collapsed" to a single-stage one (see, e.g., Savage [1954]), and enables an extensive game to be represented in a "normal" form, etc.

Yet, the fact that preferences may well change over time has been bothering economic theorists for over four decades. Allais [1947] was probably the first (in modern times) to deal with "exogenously" changing preferences, i.e., preferences which depend on time alone. "Endogenous" changes in preferences, namely those resulting from the actions of the decision-maker or other players, were studied in the early 1950's (see Schoeffler [1952] and Harsanyi [1953]).

Situations in which priorities change are plentiful; people change, events influence our perception of the world. When one is young one loves junk food, both for its taste and for the opportunity its consumption gives to meet young friends, old and new, and have a great time with them. As one gets older, one's stomach becomes more sensitive, one also

¹ Expressed by Martin Shubik (oral communication).

associates with different people, so junk food is no longer attractive. When one is young, one wants to spend a lot of time on leisure; when one gets older, providing for one's family takes priority. One might enjoy watching ceremonies of queer cults, but there is a positive probability that one may become brainwashed, desert one's family and follow the cult—a prospect that is not attractive before going to the ceremony. The reader can certainly provide many more examples of changing priorities.

It is wise to take into account the possibility of changing priorities in any long-range plan. The question then becomes how to model these situations and what strategies to recommend, as well as what strategies are likely to be played. That is the subject of this paper.

Even a brief survey of the literature is beyond the scope of this paper. A very partial list of relevant works includes Strotz [1956]; Pollak [1968]; Phelps and Pollak [1968]; Pollak [1970]; Von Weissaecker [1971]; Blackorby, Nissen, Primont, and Russell [1973]; Peleg and Yaari [1973]; Hammond [1976]; Pollak [1976]. We will briefly discuss some of these in Section 3, as a background to the presentation of our solution concept.

In Section 2 we present an extensive form model and discuss its relevance to our topic. In Section 3 we propose strategy combinations, which we call *credible equilibria*, to handle the above situations. In Section 4 we show that the set of credible equilibria is identical to the set of Nash equilibria if priorities do *not* change during the play. Thus, our solution is an extension of the Nash non-cooperative solution to situations in which priorities change during the play. We then prove, among several other results, that the set of credible equilibria contains the set of perfect equilibria of the agent-form game; hence it is never empty. Section 5 studies the set of credible equilibria. It shows, among other things, that a credible equilibrium path is also a path of a Nash equilibrium for the agent-form game. Section 6 provides some examples designed to illustrate characteristics of credible equilibria. Section 7 shows that an extension of the concept to a model in which time is a part of the data does not yield new equilibrium points. Section 8 discusses some possible variants of the concept.

Basically, we are dealing in this paper with individuals having various utilities during the

play of a game. This suggests that our model could be applied to another, very important class of situations, where a player is a *group of individuals*, a state, a party, etc. Such a player is to some extent a decision-making unit, but it does not have a utility of its own. Rather, it represents various groups, each endowed with its own utility function. For example, a state may represent farmers, manufacturers, ordinary citizens, etc., but there is no such thing as “a utility of the state”. In Section 9 we discuss the applicability of our model to such situations and show that the scope of such application is limited. Thus, an extension of our model is highly desirable.

2. The Model

We start with a *game form*² (T, P, U, C, p) . Here, T is a tree, $P = \{P_0, P_1, \dots, P_n\}$ is the *players' partition*³ of the nodes of T (P_0 is the set of Chance's nodes); $U = (U_0, U_1, \dots, U_n)$, where $U_i = \{u_{ij}\}_{j=1}^{k_i}$ is the partition of P_i into *information sets*⁴ (elements of U_0 are singletons); $C = \{C(u_{ij})\}_{i=1,2,\dots,n;j=1,2,\dots,k_i}$ is a correspondence, where $C(u_{ij})$ is the *set of choices* which are available to player i at information set u_{ij} ; $p = \{p(u_{0j})\}_{j=1,2,\dots,k_0}$ is a vector-valued function, where $p(u_{0j})$ is a probability distribution on Chance's choices at u_{0j} . For further information concerning this notation see, e.g., Selten [1975].

We assume that the game form is a *game of perfect recall* in the sense of the following definition (taken from Selten [1975]):

DEFINITION 2.1. A game form (T, P, U, C, p) is said to be of *perfect recall* if, for every i , $i = 1, 2, \dots, n$, and every two information sets u_{ij} and u_{ik} of the same player i , if one node y , $y \in u_{ik}$, comes after⁵ a choice c at u_{ij} , then *every* node x in u_{ik} comes after the same choice c .

By Kuhn's theorem (Kuhn [1953], Selten [1975]), we can and will restrict ourselves to behavioral strategies.

We shall also talk about the derived *agent-form* game, obtained by placing, for each i , different *agents* $i.j$ at the different information sets u_{ij} of player i . Each agent $i.j$ will play dual roles: on the one hand we shall regard him as a decision-making unit that acts in accordance with his own utility function. On the other hand, in reality he is the same player i located at a certain stage of the play. To complete the description of our model, we endow each agent $i.j$ with a von Neumann-Morgenstern *utility function* $h_{i,j}$, defined on lotteries over endpoints of T (which represent *pure outcomes*). Formally, therefore, our

²I.e., a game in extensive form without payments at the endpoints.

³The players are $1, 2, \dots, n$. 'Chance' is denoted by 0.

⁴To complete the description, we add that for each information set u_{ij} there are $m(u_{ij})$ edges going out from each node of u_{ij} . They are grouped into $m(u_{ij})$ disjoint equivalent classes, where each equivalent class consists of one edge from each node of u_{ij} . The equivalent classes are called the *choices*. We allow information sets with a single choice.

⁵I.e., the path from the root to x contains an arc of choice c .

game with utilities⁶ changing during the play is a six-tuple

$$(2.1) \quad \Gamma = (T, P, U, C, p, h),$$

where T, P, U, C, p are as above, and $h = (h_1, h_2, \dots, h_n)$, where, for an endpoint z ,

$$(2.2) \quad h_i(z) = (h_{i.1}(z), h_{i.2}(z), \dots, h_{i.k_i}(z)).$$

For some j 's, and in particular for *all* j 's such that $i.j$ lies on the path from the root to endpoint z , $h_{i.j}(z)$ is the utility payment of agent $i.j$ for the endpoint z of T . For other j 's $h_{i.j}(z)$ is allowed to be undefined.

DISCUSSION. The above construction comes to model a non-cooperative game in extensive form, in which the players' priorities may change during the play. To understand the relation between "reality" and our model, we provide some explanation and also discuss a possible objection to the model.

- (1) As in classical game theory, an "outcome", represented by an endpoint z of T , is the aggregate of everything that happens along the path from the root to the endpoint.
- (2) We shall be interested in this paper in certain equilibrium points which, in one application, can be viewed as possible agreements that can be reached by the players at the start of the game. For this application, *every decision must be based on what the players think at the start of the game*. Accordingly, the function $h_{i.j}$ should be interpreted as *that utility function player i believes at the start of the play he will have when he reaches information set u_{ij}* . We use here the words "believe" or "knows" in the sense of "ascribing probability one". (See the discussion in Aumann and Brandenburger [1991] concerning the relevance of this meaning.) Thus, we allow for the possibility that later on a player will find out that what he knew was wrong, in ways that he did not expect at all. When a person takes a decision, the only thing that matters is what he knows, or believes, at the moment that the decision takes place.

⁶We use the word "preferences" when we discuss the "real" situation. Their representations in the mathematical model will usually be called "utilities".

- (3) We assume that player i knows his utility function $h_{i,j}$. There is no loss of generality in this assumption. If he is not sure, being a Bayesian, he has some probability distribution over various possible utility functions. This he can represent by introducing chance moves as done in Harsanyi's theory of games with incomplete information (Harsanyi [1968]). Similarly, there is no loss of generality in assuming that all components of Γ (in (2.1)) are common knowledge.
- (4) The utility functions are merely numerical expressions for the agents' preferences. Thus, *tautologically* one expects each agent to act in accordance with maximization of his own utility function. This trivial remark is not always understood, and the misunderstanding leads to a lot of confusion. Consider the 1-person game of Figure 1.




Figure 1. The young/old me game

In this game L_1 and L_2 mean “spend on holiday vacations” and R_1 and R_2 mean “save for home for the aged”. We heard again and again that the right thing for agent 1.1 to do is to take R_1 . Indeed, agent 1.2 will certainly choose R_2 , so agent 1.1 should cooperate in order to eventually enjoy a good home for the aged (100 is a pretty large number!). This reasoning is totally wrong: Had young-me cared for old-me, this should have been reflected in 1.1's utility function. But the data shows clearly that young-me prefers to spend money on vacations rather than worry

about old age. In fact, his preference is for travel when he becomes old (see the 20), but he believes that, as an old person, his priorities will be different; therefore, he has no chance of getting the 20. Even without a theory, it should be clear that the “right” solution for this game⁷ should be L_1R_2 .

- (5) Another question that is frequently asked is this: Why do we need a new model? Is not the change of utilities simply a matter of gaining experience, or learning? To explain this, let us compare the game in Figure 1 with the one in Figure 2.




Figure 2. Taking a course in modern music

In this game, player 1 is a person who cannot stand modern music. He believes, however, that if he takes a course in modern music, he will get accustomed to this kind of music and learn to love it. L_1 , L_2 mean “stay home”. R_1 means “take the course”. R_2 means “listen to modern music”.⁸

In this case, contrary to the previous example, there is no change in the priorities of the agents. The utilities of the agents remain the same, in as much as they are

⁷For those who are not convinced, let us replace the meaning of L_1 , L_2 to be “stay home”, and R_1 , R_2 to be “consume heroin”. Agent 1.2 is already addicted to heroin. Would one still claim that a rational agent 1.1 should choose R_1 ?

⁸We put a blank as a utility for agent 1.2 after L_1 , because it makes no sense to talk about the utility of an “educated” person (agent 1.2) for the prospect of not being educated. However, we could define agent 1.2 as an agent of player 1 at a *date* when the course is over. With this interpretation, the blank could have been replaced by 10.

defined. Player 1 expects to learn how to enjoy modern music and his expectation from the course matches what he believes will be after he takes the course. This example is equivalent to the ordinary 1-person game of Figure 3.




Figure 3. A classical representation of the previous game

- (6) A possible objection to our model is this: Following logical positivism, some people feel that to talk about priorities is sheer “metaphysical nonsense” if they are not derived by observing actual *decisions*. According to this view, utilities must be derived from actual revealed preferences. If only revealed preferences count, it makes no sense to talk about revealed preferences of a future agent. How can one observe at the present time commitments to be taken 20 years from now? This is a serious criticism and it requires an honest answer:⁹

(i) This is a criticism of the whole field of game theory, not only of our model. In fact, there is hardly any application of game theory that is based on actual measurements of utilities. Game theory (including our model, we hope) is useful for the insight it sheds on real situations, for recommendations based on rough evaluation of priorities, for theoretical analysis and for clarification of issues. But we have to admit that actual measurements of utilities are usually impossible and, in those cases where they are possible, they are unreliable. In this connection see Aumann [1985].

⁹This objection was raised also in Peleg and Yaari [1973].

(ii) The belief that only revealed preferences count does not make much sense¹⁰ It can be criticized both on practical and theoretical grounds: With this belief one cannot measure, with any reasonable degree of precision, a cardinal utility such as the utility of von Neumann and Morgenstern. This is because one cannot present the players with simultaneous revealed-preference situations. What is done cannot be undone, and once a commitment is taken (if taken seriously), it cannot be cancelled.

The conceptual difficulty that we have with the revealed-preference view is this: Suppose you decide, by observing some of his actions, that a player prefers A to B and B to C . What reason do you have to believe that he would have manifested the same priorities had the choices been presented to him in a different order? Experiments that cannot be repeated are of little value! Thus, restriction to revealed preferences is not only useless, it is also questionable.

(iii) Whether we like it or not, people seldom measure utilities. They deduce their priorities (and the priorities of others) by *introspection*. To be sure, past experience takes part in this act of introspection and this includes observations of revealed preferences in similar situations in the past, but the bottom line is that important decisions are derived from priorities that result from introspection. It is these priorities that count, even if they cannot be measured by an outside observer.¹¹ Now, if we are talking about introspection, it does not matter much if we discuss a present situation or a future one. If I want to estimate my priorities, say 20 years from now, I shall observe conditions and behavior of old people, perhaps look at rudimentary available statistics about their illnesses and sufferings, and also look for things that make them happy. I shall combine these facts (consciously or subconsciously) with what I think I know about myself and—rightly or wrongly—deduce my future priorities by introspection. It is the outcome of my introspection that will dictate my decisions.

To sum up: Attempts to determine von Neumann-Morgenstern utilities by observing people's behavior did not prove successful. In fact, experts on experimental economics¹²

¹⁰We employ here the tactic that an attack is the best defense...

¹¹Savage already recognized their importance when he discussed "conceivable acts" (Savage [1972]).

¹²Oral communication with Reinhard Selten and Martin Shubik.

claim that real people do not seem to behave as utility maximizers. One can take a view, as expressed by John Harsanyi,¹³ that game theory is a theory about ideal people—people who do not exist in the real world, and enjoy the subject for its aesthetic value. *In this sense our model is certainly as useful as the classical one.* And we can be more flexible, recognize that decisions are made by the process of introspection, the results of which govern the players' actions. Utility theory and game theory can help the decision-maker by leading his thoughts to the required priority determinations, thereby “educating” him to be more systematic and logical, hoping that with this help his decisions will be reached faster and with fewer “oversights”. At any rate, *both conceptually and practically*, introspections about future prospects are not more complicated than introspections about the present.

All the above has very little to do with another main application of game theory; namely, to shed more light on conflict situations, to deepen our understanding of the conflict per se, even though in reality most conflicts cannot be quantified. In this respect, our model is certainly an interesting and useful extension of the classical one.

¹³Oral communication.

3. Credible Equilibria

The purpose of this paper is to generalize the concept of Nash equilibrium to games with utilities changing during the play.

One idea that comes to mind is to recommend *an agent-form Nash equilibrium*. It seems that such an equilibrium is satisfactory; no agent will deviate because, by deviation alone, he has nothing to gain. Indeed, this recommendation was studied in the important paper of Peleg and Yaari [1973]. Nevertheless we now feel that this recommendation *is not good enough*. The reason is that Nash equilibrium is robust with respect to a deviation of a single agent—not against deviation of several agents. But several agents of the same player can cooperate quite easily, to the benefit of all of them, *because they are all the same individual*. Thus, an agent-form Nash equilibrium is not necessarily stable. It seems that what we need is some kind of coalition-proof equilibrium,¹⁴ where coalitions are restricted to agents of the same player. Unfortunately, if we extend that definition directly, we may easily reach situations in which no such equilibrium exists. Moreover, we may unjustifiably reject “good” points: Suppose a strategy combination is rejected because a certain coalition of agents of a player can deviate and do better—it may well be that this deviation will not be obeyed, because some agents of the same player (not necessarily a subset of the deviating players) can do better after the deviation is “adopted” and cause loss to some members of the deviating coalition. If this is the case, the original strategy combination seems to us quite reasonable.

To overcome the above difficulties, we introduce here the concept “*credible deviation*”, defined recursively, and define “*credible equilibrium*” as one at which no credible deviations exist.

To make things precise, we need to establish some notation. Let $\Gamma = (T, P, U, C, p, h)$ be a game of perfect recall with utilities changing during the play. A behavioral strategy $s_{i,j}$ of agent i,j is a probability distribution over the choices c_{ij} at u_{ij} . We note by $S_{i,j}$ the set of these strategies. A behavioral strategy for player i is the k_i -tuple $s_i := (s_{i,1}, s_{i,2}, \dots, s_{i,k_i})$,

¹⁴A concept that was introduced in Bernheim, Peleg and Whinston [1987] and in Peleg [1992].

where $s_{i,j}$ belongs to $S_{i,j}$. Thus, the set of behavioral strategies for player i is $\times_{j=1}^{k_i} S_{i,j} =: S_i$. An n -tuple of behavioral strategies is $s = (s_1, \dots, s_n)$ and the set of these n -tuples, called also “points” is $S = \times_{i=1}^n S_i$.

Let Q be a set of agents (to be thought of as belonging to the same player). We denote by $-Q$ the set $M \setminus Q$, where M is the set of all agents (not only of the same player). For a point s , we denote by s_Q the vector of strategies¹⁵ $(s_{i,j})_{i,j \in Q}$. Similarly, $S_Q := \times_{i,j \in Q} S_{i,j}$. For simplicity we also write

$(s_{-i,j}, s'_{i,j})$ which expresses a deviation of agent i,j to $s'_{i,j}$, instead of the more precise notation $(s_{-\{i,j\}}, s'_{i,j})$.

Given an n -tuple of behavioral strategies s , it induces a probability distribution on the endpoints. Thus, we can denote by $h_{i,j}(s)$ the utility of agent i,j for this lottery. For s and s' in S , we write $s' \succ_{i,j} s$ iff $h_{i,j}(s') > h_{i,j}(s)$.

Finally, for a given s and a given agent i,j , we denote by $\Gamma_{i,j}^s$ the game obtained from Γ by converting the strategy of every agent, other than the strategies of i,j and the agents of player i that play *after him*,¹⁶ to chance mechanisms playing as in s (and converting the information sets of the chance mechanisms to singletons). $\Gamma_{i,j}^s$ is the game that the players think at the start of the game that agent i,j is facing, given that every agent other than he and his followers follow s . Note that $\Gamma_{i,j}^s$ is a 1-person game—possibly with several agents.

We can now introduce our solution concept. Note that, because the game is of perfect recall, the probability distributions on the nodes of the information sets of agents of player i , given that they are reached, depend only on the strategy $(n-1)$ -tuple¹⁷ s_{-i} and not on s_i . Indeed, if a non-singleton information set of an agent i,j_0 is reached, all previous choices taken by player i are known, the information set results from not knowing choices taken by chance and by other players. The probabilities on its nodes can be calculated from p and s_{-i} .

DEFINITION 3.1. Let $\Gamma = (T, P, U, C, p, h)$ be a game of perfect recall with utilities changing during the play. Let s be an n -tuple of behavioral strategies. Let Q be a set of agents

¹⁵Strictly speaking, we should fix an order on the agents to make it a vector.

¹⁶We say that i,j plays after i,j_0 if every path from $u_{i,j}$ to the root passes through u_{i,j_0} .

¹⁷ s_{-i} is a short notation for $(s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$.

of player i , $i \neq 0$, containing an agent $i.j_0$ and some of i 's agents that play that play after $i.j_0$. A strategy $|Q|$ -tuple s'_Q is said to be a *credible deviation from s , struck by agent $i.j_0$ using Q* , if:

- (i) $s' \succ_{i.j_0} s$, where $s' := (s'_Q, s_{-Q})$, i.e., agent $i.j_0$ strictly prefers that everyone play according to s' rather than everyone play according to s .
- (ii) $s' \succ_{i.j} (s_{-i.j}, s'_{i.j})$ for all $i.j \in Q$, $i.j \neq i.j_0$. Thus, agent $i.j$ strictly prefers to play $s'_{i.j}$ rather than $s_{i.j}$, given that the other agents play as dictated by s' . We shall refer to this condition by saying that $i.j$ prefers to *comply* with s' .
- (iii) No agent of i , whether in Q or not, that plays *after* $i.j_0$, can strike a credible deviation from s' .

Note that the definition is not circular, because the relation “plays after” is acyclic.

DISCUSSION. We view the deviation as a set of instructions s'_Q given by $i.j_0$ to the members of Q . Condition (i) states that $i.j_0$ prefers that these instructions are obeyed, given that agents of $-Q$ continue¹⁸ to follow s . Indeed, otherwise, why did he give such instructions? Note that this implies that $i.j_0$ is reached with positive probability under s , because prior to $i.j_0$ there is no distinction between elements of s and s' . Condition (ii) implies that each member $i.j$ of Q is reached with positive probability under s' . When $i.j$ comes to play he has two suggested strategies: The original suggestion $s_{i.j}$ dictated by s and the deviation suggestion $s'_{i.j}$. Which one will he obey? The condition states that each member of Q , except¹⁹ perhaps $i.j_0$, actually prefers to comply with s' . In other words, when $i.j$ of Q comes to play (equivalently, when $i.j$ considers playing in $\Gamma_{i.j}^{s'}$), he prefers $s'_{i.j}$ to $s_{i.j}$, given that other agents follow s' . Thus, such an agent has an incentive to follow the instructions. Following the spirit of Nash equilibrium, we took the position that even if $i.j$ is indifferent between $s'_{i.j}$ and $s_{i.j}$, he will not switch²⁰ to $s'_{i.j}$.

But there remains some doubt: What guarantees do the agents have that every agent of i will follow s' ? Perhaps one can strike a deviation from s' ! Condition (iii) is introduced

¹⁸Subsequently we shall see that he can count on this.

¹⁹This exception was put in order to allow for $i.j_0$ to continue using $s_{i.j_0}$. If other agents are asked to continue using $s_{i.j}$, we simply do not include them in Q .

²⁰One could think of a different position (see Section 8).

to prohibit this possibility. The players that come to play after $i.j_0$ can rest assured that none of them can deviate from s' in a credible way. Condition (iii) gives our definition a recursive flavor, and it works because, for every agent $i.j$ after $i.j_0$, the “ i -length”²¹ of $\Gamma_{i,j}^{s'}$ is smaller than the i -length of $\Gamma_{i,j_0}^{s'}$.

REMARK. Although it is reasonable to assume that an agent of i not only remembers what choices other agents of i took in the past, but also knows by what probability distribution they were obtained, we *do not* make this assumption. So, if, say, an agent $i.j_1$, decides to “cheat” and move from a completely mixed strategy s_{i,j_1} to another one, any agent $i.j_2$ after him will not recognize that cheating took place. As has been said above, knowing s_{-i} and knowing that he was reached, is sufficient, in a game of perfect recall, for agent $i.j_2$ to compute all the probabilities at his information set. In fact, sometimes an agent $i.j$ can predict that $i.j_0$ was cheating and this will not change his evaluation and behavior, as the next example shows.

EXAMPLE 3.2. Consider the 1-player 3-agent game of Figure 4.




Figure 4. A credible deviation that will be violated

²¹By “ i -length” we mean the longest path from an agent $i.j$, $i \neq 0$, to an endpoint.

Let $s_{1.1} = (1/2, 1/2)$, $s_{1.2} = (1, 0, 0)$, $s_{1.3} = (1, 0, 0)$. A credible deviation by 1.1, using all the agents, may be: $s'_{1.1} = (1/2, 1/2)$, $s'_{1.2} = (0, 1, 0)$, $s'_{1.3} = (0, 1, 0)$. It is credible, because *after* 1.1 plays, every agent is gaining and maximizing under s' . Nevertheless, 1.1 is actually likely to deviate and play $(0, 1)$, to get 3 instead of 2.5. Now, if 1.2 knew that he was cheating he could punish him for his “betrayal”, say, by playing $(0, 0, 1)$, but we do not need to go into such considerations. In fact, we shall later prove (Theorem 5.1) that if $i.j_0$ has a credible deviation from s , then he has another credible deviation from s in which it does not pay him to move away.

With the above discussion and remark we can now state our main definition.

DEFINITION 3.3. Let Γ be a game with perfect recall and utilities changing during the play. An n -tuple s of behavioral strategies is called a *credible equilibrium* if there are no credible deviations from it; i.e., if there does not exist an agent $i.j_0$ and a coalition Q and a vector of behavioral strategies s'_Q which constitute a credible deviation from s .

4. Existence theorems

We start by showing that in ordinary games, the set of Nash equilibria (Nash [1951]) and the set of credible equilibria coincide. This substantiates the claim that we are extending the Nash's solution to games with utilities changing during the play.

THEOREM 4.1. *Let $\Gamma = (T, P, U, C, p, h)$ be a game of perfect recall in which, for each player, the utility functions of his agents coincide.²² Let Γ^* be the ordinary representation of this game²³. The set of credible equilibria (CrE) of Γ is equal to the set of Nash equilibria (NE) of Γ^* .*

PROOF: If $s \in \text{NE}$ in Γ^* then $s \in \text{CrE}$ in Γ , because no deviation by an agent of a player, using a set of agents of that player, can satisfy condition (i) of Definition 3.1.

Conversely, if $s \notin \text{NE}$ in Γ^* , then there exists a deviation s'_Q by player i , which is a best reply to s_{-i} and player i , and therefore each of his agents, prefers $s' := (s'_Q, s_{-Q})$ to s . Let $i.j_1, i.j_2, \dots, i.j_k$ be the “first members” of Q ; namely, the members of Q whose paths to the root do not contain other members of Q . Because the game is of perfect recall, each agent $i.j_\nu$, $\nu = 1, 2, \dots, k$, can compute his expected payoff in $\Gamma_{i.j_\nu}^s$, both under s and under²⁴ s' . The difference between these expected payoffs depends only on actions taken by the members of $Q_\nu := \{\text{members of } Q \text{ who are agents in } \Gamma_{i.j_\nu}^s\}$. The strategy combination s'_{Q_ν} , which is s' restricted to Q_ν , is a maximizing strategy of player i in $\Gamma_{i.j_\nu}^s$. Moreover, since $s \notin \text{NE}$, at least one agent, say $i.j_1$, is reached with positive probability under s and strictly prefers $s' := (s'_{Q_1}, s_{-Q_1})$ to s in $\Gamma_{i.j_1}^s$. We now modify s'_{Q_1} by working backwards from the endpoints on members of Q_1 . If an agent $i.j$ in Q_1 , in his turn, is indifferent between $s'_{i,j}$ and $s_{i,j}$, *given that his information set is reached*²⁵ and the strategies of players after him have been already determined, let him switch to $s_{i,j}$.

²²Strictly speaking, we should have added “whenever defined” (see Figure 2). However, credible equilibria do not depend on utilities of an agent off paths in which he plays, so we might as well assume that the utility functions are defined at all endpoints.

²³Compare Figures 2 and 3.

²⁴I.e., when his followers in $\Gamma_{i.j_\nu}^s$ play as dictated either by s , or by s' , respectively.

²⁵We also assume that he *can be reached* with positive probability if player i plays appropriately. Remember that in this case the probability distribution on the nodes of his information set is determined by the strategy combination s_{-i} . If he cannot be reached, no matter what i does, he cannot compute the probabilities, but we can safely require that he plays $s_{i,j}$.

After each switch, player i will still be playing a best reply in Γ^* . After all these switches there will be agents of Q_1 that are reached with zero-probability under the new strategies. Let them too switch to play as in s (if they have not already switched). Denote by Q_2 the members of Q_1 who still play as in s' , $Q_2 \neq \emptyset$. s'_{Q_2} is still a maximizing strategy in Γ_{i,j_1}^s . We claim that s'_{Q_2} is a credible deviation from s , struck by $i.j_1$. Indeed, we already know that $i.j_1$ prefers $s' := (s'_{Q_2}, s_{-Q_2})$. Q_2 consists of exactly those agents who prefer to follow s' when they come to play. Finally, all agents of i have the same utility function and, after $i.j_1$, no agent can do better than the above maximizing strategy; so, no agent after $i.j_1$ can strike a credible deviation from s' . ■

It follows from Theorem 4.1 that CrE is not empty for ordinary games of perfect recall. However, we need to establish existence in general as well. This will be done as a corollary to the following theorem which shows that CrE contains the set of *agent form perfect equilibria* (APE) (see Selten [1975]).

THEOREM 4.2. *Let $\Gamma = (T, P, U, C, p, h)$ be a game of perfect recall with utilities changing during the play. Let Γ^* be the agent form game obtained from Γ by considering different agents as different players. If s is a perfect equilibrium in Γ^* then s is credible in Γ .*

PROOF: s is a limit of a “test sequence” (s^k) , $k = 1, 2, \dots$, where $s^k \in \text{NE}$ in a “perturbed agent form game” Γ^k . Γ^k is a game having the same T, P, U, C, p, h as Γ , but the behavioral strategies at each choice c are restricted so that the probability to choose c is not smaller than some positive number δ_c^k , with $\sum_{c \in C(u_{ij})} \delta_c^k < 1$ for every information set u_{ij} , and $\lim_{k \rightarrow \infty} \delta_c^k = 0$, all c .

Suppose s is not credible in Γ , then there exists an agent $i.j_0$ who can strike a credible deviation s'_Q . Let $i.\ell$ be a *last agent*²⁶ in Q . Agent $i.\ell$ prefers to play $s'_{i,\ell}$ rather than $s_{i,\ell}$ in $\Gamma_{i,\ell}^s$, given that the agents that follow him play as in s , because for these agents there is no distinction between s' and s and because the deviation was credible. This preference will not change if we modify $s'_{i,\ell}$ and s slightly. Modify $s'_{i,\ell}$ to $s''_{i,\ell}$ which is positive in all components, so that it is a legitimate strategy in Γ^k for sufficiently large k . Modify s to s^k and one obtains $s''_{i,\ell} \succ_{i,\ell} s^k_{i,\ell}$ in $(\Gamma^k)_{i,\ell}^{s''}$ for sufficiently large k . This means that s^k is not in

²⁶Namely, an agent that after him all agents, if exist, play in s' as in s .

NE in Γ^k , because in Γ^k $u_{i,t}$ is reached with positive probability under s^k , a contradiction.

■

COROLLARY 4.3. *The set of credible equilibria is never empty.*

PROOF: Theorem 4.2 and the fact APE is not empty (Selten [1975]). ■





Figure 5. A credible equilibrium not in APE

One may wonder if all credible equilibrium points are in APE. That this is not the case even in ordinary games follows from Theorem 4.1, because there are ordinary games with $NE \neq APE$. For a simple example which is not an ordinary game consider Figure 5. It shows that the concept of CrE permits a certain amount of punishment.

REMARK. An agent form subgame perfectness is not sufficient for credibility. A simple example is given in Figure 6. (The arcs denote the credible deviation.)

Figure 6. An agent form subgame perfect equilibrium which is not in CrE



5. Some properties of the set of credible equilibria

In this section we shall look more critically at the concept “credible equilibrium”. We shall try to find flaws and investigate their seriousness. Take, first, the concept “credible deviation”. Suppose s is not credible. Can we always recommend a “best credible deviation”? The answer is negative as can be seen in Figure 7.




Figure 7. A case in which there does not exist a best credible deviation

Here, $s = (R_1, R_2, L_3)$ is not credible. The only credible deviations are those struck by agent 1.1, instructing all the agents to move to $s' := (L_1, L_2, [(\frac{1}{2} - \varepsilon)M_3, (\frac{1}{2} + \varepsilon)R_3])$, $0 < \varepsilon < \frac{1}{2}$. Note that s' is a credible equilibrium. Of course, agent 1.1 would like to choose ε as small as possible, but he cannot take $\varepsilon = 0$, because then 1.2 will not cooperate. Note that this is a consequence of the position we took in Definition 3.1, that cooperation can only be taken if there is real gain.

A more serious criticism, at least aesthetically, is the fact that CrE is not necessarily closed. Consider, for example, the 2-person game in Figure 8, in which agents 1.1 and 1.2 move L_1 and L_2 , and agent 2.1 mixes between L_3 and R_3 with probabilities $(1 - \varepsilon, \varepsilon)$. For

every $\varepsilon, 0 < \varepsilon \leq 1$, this is a credible equilibrium. It ceases to be, if $\varepsilon = 0$.




Figure 8. A case in which CrE is not a closed set

Note that this example hinges on our conservative stand that an agent will cooperate only in face of a gain.

In Section 3 we presented an example of a credible deviation (Example 3.2) that was certain to be violated by the deviating agent himself. This was possible because of the need to keep Definition 3.1 recursive, and therefore meaningful. We argued that nevertheless each agent of Q , after $i.j_0$, would still prefer to obey s'_Q . Fortunately, we do not have to defend this argument any further, in view of the following theorem which shows that we may just as well restrict s'_Q to cases in which $i.j_0$ too cannot further violate.

THEOREM 5.1. *If s'_Q is a credible deviation from s by an agent $i.j_0$, then there exists a deviation s^*_G from s , by the same agent, such that he too cannot strike a credible deviation from $s^* := (s^*_G, s_{-G})$.*

PROOF: The proof involves several steps.

STEP 1. Denote by A' the set of pure choices used by $i.j_0$ with positive probability under $s'_{i.j_0}$. We can and do assume that by playing $s'_{i.j_0}$ agent $i.j_0$ is maximizing under the condition that he must choose from A' and that all the agents that play after him (not necessarily only agents of i) obey s' . Indeed, if not, replace $s'_{i.j_0}$ by $\tilde{s}'_{i.j_0}$ in which $i.j_0$ is maximizing under the restrictions stated above. Let \tilde{Q} be a subset of Q , reached

with positive probability under $\tilde{s} := (\tilde{s}_{i,j_0}, s'_{Q \setminus \{i,j_0\}}, s_{-Q})$ and not containing i,j_0 . The combination $(\tilde{s}'_{i,j_0}, s'_Q)$ is a credible deviation from s , struck by i,j_0 , which satisfies the maximizing requirement.

STEP 2. Denote by B the set of all agents of i that play after i,j_0 and are reached with positive probability under s' . Denote by C the set of all agents of i that follow pathwise agents of B , but reached with probability 0 under s' .

We shall show that if there exists a credible deviation s''_D from s' , struck by i,j_0 , then $D \cap (B \cup C) = \emptyset$. Suppose $D \cap B \neq \emptyset$ then there is a last agent i,j_1 in $D \cap B$, i.e., an agent such that those that follow him, and play differently under s' and s'' , are reached with probability 0 under s' . Let D^* consist of i,j_1 and the agents in D that play after him. We claim that s''_{D^*} is a credible deviation from s' , struck by i,j_1 —thus arriving at a contradiction because s' was credible. To verify the claim, observe that conditions (ii) and (iii) of Definition 3.1 are satisfied for s''_{D^*} , because s''_D was a credible deviation struck by i,j_0 and because every player after i,j_1 is also a player after i,j_0 . More care is needed to verify condition (i). Consider the game $\Gamma_{i,j_1}^{s'}$ and denote by \tilde{s}'' the restriction of s'' to this game. We know that in $\Gamma_{i,j_1}^{s'}$, $\tilde{s}'' \succ_{i,j_1} (\tilde{s}''_{-\{i,j_1\}}, s'_{i,j_1})$, because \tilde{s}''_D was a credible deviation from s' , struck by i,j_0 and i,j_1 plays after him (and is reached with positive probability both under s' and under s''). What we have to show is that in $\Gamma_{i,j_1}^{s'}$, $\tilde{s}'' \succ_{i,j_1} (s'_{-\{i,j_1\}}, s'_{i,j_1})$. This, in fact, is the case, because i,j_1 was a last agent in $D \cap B$. Thus, in $\Gamma_{i,j_1}^{s'}$, every agent i,j after i,j_1 , for whom $s''_{i,j} \neq s'_{i,j}$, is reached under s' with probability 0, so there is no change in payoffs if we require such agents to play $s''_{i,j}$ instead of $s'_{i,j}$.

Since $s''_{i,j} = s'_{i,j}$ for every agent in B , it follows that every agent in C is reached with probability 0 under s'' , therefore $D \cap C = \emptyset$.

STEP 3. We have shown that $s'' = s'$ on paths after i,j_0 , reached with positive probability under s' and their continuations. Thus, the support A'' of s''_{i,j_0} must contain choices other than the choices in A' , then $A'' \cap A' = \emptyset$ and s'' directs the play to paths disjoint from those reached with positive probability under s' and their continuation. As in Step 1, we can and do assume that i,j_0 is maximizing given that he is restricted to A'' and the other agents play according to s'' .

STEP 4. If there exists a credible deviation s'''_E from s'' , struck by $i.j_0$, then, as in Step 2, $E \setminus \{i.j_0\}$ does not contain agents of i reached under s'' with positive probability or agents of i that follow such agents. By arguments similar to those given in Step 2, it does not contain agents of i reached with positive probability under s' , and pathwise followers of such agents. Thus, if $s'''_{i.j_0}$ is maximizing on its support, it directs the play to paths different from the paths supported by s' or s'' .

STEP 5. We continue in this fashion until we reach a credible deviation $s^{(k)}_G$ from an $s^{(k-1)}_F$, struck by $i.j_0$ and player $i.j_0$ cannot strike a credible deviation from $s^{(k)}_G$, neither alone, restricted to the support of $s^{(k)}_{i.j_0}$, nor by directing the play to other paths, because such paths are not available.

STEP 6. $s^{(k)}_G$ is also a credible deviation from s , struck by $i.j_0$. Indeed by transitivity, $s^{(k)} \succ_{i.j_0} s^{(k-1)} \succ_{i.j_0} \dots \succ_{i.j_0} s' \succ_{i.j_0} s$; for every agent $i.j$ of G , $s^{(k-1)}_{i,j}$ is identical to $s_{i,j}$ and he prefers to comply with $s^{(k)}$; and neither $i.j_0$, nor his followers can strike a credible deviation. ■

A credible equilibrium need not be ANE, because agents off the paths of the play may sometimes act “irrationally”. Figure 9 provides such an example. However, if one thinks this is undesirable, one can find comfort in the following theorem.




Figure 9. A credible equilibrium which is not in ANE

THEOREM 5.2. *The paths that are supported by a credible equilibrium are also supported by a credible equilibrium which is also a Nash equilibrium for the agent form game.*

For the proof we need the following

LEMMA 5.3. *Let s be a credible equilibrium and let i be a player. There exists an n -tuple of strategies \tilde{s} , such that*

- (1) \tilde{s} induces the same probability distribution as s on the set of nodes. In fact, it coincides with s at all information sets reached under s with positive probability.
- (2) Every agent $i.k$ of i is maximizing in \tilde{s} , namely he is maximizing his expected payoff given that all other agents obey \tilde{s} .
- (3) \tilde{s} is a credible equilibrium.
- (4) Every agent $j.k$, $j \neq i$, if he was maximizing under s , he is still maximizing under \tilde{s} .

PROOF: As usual, we consider s_{-i} fixed and known to player i . We construct \tilde{s} successively, backwards eliminating agents $i.k$ who are not maximizing.

Let $i.k$ be a last agent who is not maximizing. Temporarily replace $s_{i.k}$ by a pure strategy $\hat{s}_{i.k}$ which is a best reply for him. This necessarily directs the play to paths not supported by s , because s was a credible equilibrium. Along these new paths we modify s by backward induction on the agents of i , after $i.k$, instructing each of them to employ best reply, given that he is reached,²⁷ and keeping with s whenever there is indifference and whenever he cannot be reached. Agent $i.k$ cannot gain by these modifications, compared to s , because if he could, he could also strike a credible deviation from s , giving himself and his followers the same instructions. So we let him revert to $s_{i.k}$. The resulting modification satisfies (1) with \tilde{s} being the modified strategy n -tuple. It is also credible, because all modifications were done on unreachable paths that only $i.k$ could direct to them, which he would not because it were not profitable.

Agent $i.k$ may perhaps, still gain by picking up another pure strategy instead of $\hat{s}_{i.k}$. We

²⁷ We remind the reader that if an agent is reached, then, because the game is of perfect recall, he can compute the probabilities of reaching each node of his information set and these probabilities depend only on s_{-i} .

eliminate such avenues successively in a similar fashion until eventually $i.k$ is maximizing with respect to the modified strategy n-tuple.

This process is repeated until all agents of i are maximizing and so (1)-(3) are satisfied.

For agents of other players, if they are off the support of s (which is also the support of \tilde{s}), they are maximizing automatically. If they are in the support and maximizing under s , they are also maximizing under \tilde{s} , because $s = \tilde{s}$ in the support and they cannot lead the game to paths where $s \neq \tilde{s}$. This is condition (4). ■

PROOF OF THEOREM 5.2: Apply the lemma consecutively for all the players. The final modification will be in ANE, because every agent will be maximizing. Moreover, this modification coincides with s on the paths supported by s . ■

6. Examples

The purpose of the examples in this section is to illustrate some features of the credible deviations and get a better understanding of them.

EXAMPLE 6.1. Strategy (L_1, L_2) in Figure 10 is not credible in spite of the fact that agent 1.2 receives less than the original payment under the credible deviation (R_1, R_2) . Here we see the importance of the timing of decisions. When 1.2 plays, he has no alternative but to comply with the deviation. Thus, within the agents of a given player, the concept of credible equilibrium has some flavor of perfection, although in general, as we have seen, CrE may even contain points not in ANE.




Figure 10. A credible deviation that harms an agent

EXAMPLE 6.2. Strategy (L_1, L_2) is credible in the game of Figure 11. An instruction (R_1, R_2) by agent 1.1, although promising to increase agent 1.2's payoff, is not safe for agent 1.1. When agent 1.2 comes to play he has no motivation to move to R_2 . In Section 8 we shall propose other variants under which (R_1, R_2) will be considered a credible deviation.

EXAMPLE 6.3. Consider the game in Figure 12. Here, Wife (player 1) and Husband (player 2) can either "save" (s) or "consume" (c) at various stages. If both players save at




Figure 11. A credible equilibrium which is not Pareto optimal




Figure 12. The saving/consuming game

all stages, they can buy a high-quality car. Partial saving $((c_1, s_2, s_3)$ in the figure) allows them to buy only a low-quality car. Wife is indifferent between high- and low-quality car. Husband strongly prefers the high-quality car (for instance, he needs it to show off). In

fact, he would “consume” rather than buy the low-quality car. Wife can consume a small portion without Husband noticing. Alas, when consuming she enjoys her purchase all the more, and her desire to further consume becomes stronger (say, she bought an expensive ring and now she needs a necklace to match). This is evident when one compares the utility differences between 1.1 and 1.2. The rest of the story can be read from the figure.

Clearly, (s_1, s_2, s_3) is not credible: Wife 1.1 switches to c_1 . Similarly, (c_1, s_2, s_3) is not credible: Husband will spend his money on leisure. Thus, there is no chance of getting any car using pure strategies. The point (c_1, c_2, c_3) is credible. It yields $(3, 4, 5)$ to the agents. What a sad prospect. But there exists a credible behavioral strategy that allows for complete saving with positive probability. It is $([2/3(s_1), 1/3(c_1)], [1/2(s_2), 1/2(c_2)], c_3)$. Under this strategy they will buy a high-quality car with probability $1/3$ and a low-quality car with probability $1/6$. The payoff to 1.1 will be 3 and it will be 4 to 2.1. Agent 1.2 will be “created” with probability $1/3$ and will receive 5. It is interesting to note that if 1.2 had the same utility function as 1.1, the only equilibrium point would have been (c_1, c_2, c_3) .

7. Timewise Credible Equilibria

One may question the definition of a credible deviation in two ways:

- (1) Is it not an artificial requirement to allow each agent control over only one information set? What if the same agent controls several?
- (2) Why do we allow a deviating agent to instruct only agents that follow him *pathwise*? After all he may derive payments also at paths in which he does not play.

To cope with the above criticism we have to extend the model of Section 2 by introducing *dates* at each information set. This we shall do in this section and then prove that we do not get other equilibrium points. For the sake of brevity we shall omit the formal definitions, which are straightforward extensions of the previous sections. First, we shall provide two examples to give the reader some insight into the problem. In this section we use the convention that nodes in a figure on the same level are played, if reached, at the same time and that lower nodes are played, if reached, at later dates.




Figure 13. A “counterexample”

EXAMPLE 7.1. The tree in Figure 13 refers to a one-player one-agent game, but the agent occupies two information singletons A and B . It seems that this example supports the

above criticism: The player cannot benefit by deviating from (L_1, L_2) at any one node, but he can benefit by playing (R_1, R_2) .

This example is not valid because an agent cannot have two different utilities at any particular node. This is simply meaningless. We must regard A and B as occupied by different agents, even if they are assigned the same date.




Figure 14. A correct representation

EXAMPLE 7.2. Figure 14 is a correct version. It seems, however, that we get a similar counter example by allowing agent 1.1 instruct himself and 1.2 to move right.

Yes, but why would agent 1.2 obey? If he has to play he knows that 1.1 is not playing, so he knows he cannot gain by deviation.

We now extend the model of Section 2 by adding a date at each information set and by allowing the same agent control several information sets that are assigned the same date. We remind the reader that an agent has exactly one utility function at all endpoints at which it is defined. We call the various information sets of an agent “*subagents*” of that agent. We assume that all subagents of an agent are associated with the same date. Thus, an agent is a unit that makes a decision on a certain date, based on one utility function. We modify the definition of a credible deviation of Section 3 by allowing members of Q other

than $i.j_0$ to be agents of i who play not earlier *in time* than $i.j_0$. Condition (i) remains that agent $i.j_0$ prefers s' to s . Condition (ii) says that every agent $i.j$ in Q , other than $i.j_0$, prefers, *when he comes to play*, s' to $(s'_{-i,j}, s_{i,j})$. Note that when he comes to play he knows at which information set he is located, so, as far as he is concerned, only strategies by subagents that are pathwise after this information set count for this preference. Condition (iii) says that every agent after $i.j_0$, *timewise*, cannot strike a credible deviation. Since the number of time periods is finite, we have here a well defined recursive definition. We call it a “*timewise credible deviation from s , by $i.j_0$, using Q* ”. The following two lemmas and theorem prove that with this extension we do not get different credible equilibrium points.²⁸

In this section we are using the adverbs “*pathwise*” and “*timewise*”. “Pathwise” refers to notions related to the tree. “Timewise” refers to notions related to the dates. Thus, saying “playing timewise after an agent” is not the same as saying “playing pathwise after an agent”. It is reasonable to require, and we do require, that the latter implies the first.²⁹

LEMMA 7.3. *If s'_Q is a timewise credible deviation from s by an agent $i.j_0$, then there exists a timewise credible deviation from s by a subagent of $i.j_0$, in the game Γ^* , obtained from the original game Γ by regarding each subagent as a different agent.*

PROOF: Denote by $i.j_{01}, i.j_{02}, \dots, i.j_{0k}$, $k \geq 1$, the subagents of $i.j_0$. They are all supposed to choose a move on the same date, if the play reaches them. They all have the same utility function.

We regard s_{-i} as fixed and known to player i . Because the game is of perfect recall, each subagent knows the probability distribution at the nodes of his information set, if the play reaches him. Given either s_i or s'_i , he can compute his expected payoff at the endpoints.³⁰ In particular, he can compute the sum of the expected payoffs at endpoints that follow him. We shall call this sum the *pathwise expected payoff* of the subagent. The

²⁸To simplify the exposition, we assume that agents do not employ correlated strategies among their subagents. This is legitimate, because the game is of perfect recall, so only induced probability distributions at the various information sets matter.

²⁹See the Remark at the end of this section.

³⁰Note that prior in time to $i.j_0$, s and s' coincide.

expected payoff of agent $i.j_0$ is a weighted average of the pathwise expected payoffs of the subagents $i.j_{0\nu}$, $\nu = 1, 2, \dots, k$. Since $s' \succ_{i.j_0} s$, at least one subagent, say $i.j_{01}$, has a higher pathwise expected payoff under s' than under s . We shall now construct a *timewise* credible deviation from s in Γ^* , struck by agent $i.j_{01}$. It is simply s'_Q . In other words, instead of saying that $i.j_0$ struck s'_Q , we now say that $i.j_{01}$ strikes a deviation by giving the same instructions. Indeed, condition (i) is fulfilled because of the way $i.j_{01}$ was selected. When an agent $i.j_p$ in Q comes to play he knows that he is, say, at $u(i.j_{p1})$. He prefers to comply with s' in Γ^* , because he had the same preferences in Γ and his pathwise expected payoff is the same in both games. This is condition (ii). If in Γ^* there is a subagent $i.j_q$ who plays timewise after $i.j_{01}$, who can strike a timewise credible deviation from s' , then, in Γ , $i.j$ could strike the same credible deviation from s' , which is impossible. This proves condition (iii). ■

LEMMA 7.4. *Let Γ be a game in which every agent occupies a single information set. If an agent $i.j_0$ can strike a timewise credible deviation s'_Q from s , then there is an agent who can strike a pathwise credible deviation from s .*

PROOF: If there exists an agent $i.j_1$ in Q , $j_1 \neq j_0$, not pathwise after $i.j_0$, then there exists an agent $i.j_2$ in Q , not pathwise after $i.j_0$, who is *timewise last* to be reached with positive probability also under s . This agent, when he comes to play, prefers s' to $(s'_{-i.j_2}, s_{i.j_2})$, which means that he also prefers s' to s , because, by choosing $s'_{i.j_2}$ he diverts the play to paths in which there is no distinction between s and s' , or to paths in which, under s , it does not matter whether $s_{i,j}$ or $s'_{i,j}$ is played. This is condition (i) for a deviation $s''_{Q^*} := s'_{Q^*}$, struck by $i.j_2$ from s , where Q^* consists of $i.j_2$ and members of Q that play *pathwise* after $i.j_2$. Any such member other than $i.j_2$ prefers to comply with s'' , because he preferred to comply with s' and his preferences do not depend on moves taken not pathwise after him. This is condition (ii) of Definition 3.1. Similarly, no agent pathwise after $i.j_2$ can strike a credible deviation from s'' , because if he could, he could also strike the same deviation from s' . ■

THEOREM 7.5. *Let Γ be a game with dates and with agents controlling several information sets. Let Γ^* be the game obtained from Γ by regarding subagents as different agents and*

by removing the date labels. The set of timewise credible equilibria in Γ is identical to the set of (pathwise) credible equilibria in Γ^* .


PROOF: From Lemmas 7.3 and 7.4, it follows that if s is a pathwise credible equilibrium in Γ^* then it is also timewise credible equilibrium in Γ . Suppose s is a timewise credible equilibrium in Γ then, in particular there does not exist a timewise credible deviation s'_Q by a subagent $i.j_{01}$ that uses only members of Q who follow him pathwise. Therefore, for any such attempted deviation, either $i.j_{01}$ does not prefer s' to s or there exists a member of Q (pathwise) after $i.j_{01}$ who would rather comply alone with s , or an agent $i.j_1$ in Γ who plays timewise after $i.j_{01}$, can strike a timewise credible deviation s''_D from s' . In the first two cases, s'_Q is not a credible deviation in Γ^* . In the third case s'_Q would still not be a credible deviation in Γ^* if $i.j_1$'s deviation from s' uses only subagents that follow $i.j_{01}$ pathwise. If this is not the case then there exists in D a last subagent, reachable with positive probability also under s who can strike a timewise credible deviation³¹ to s . This, however, is impossible, because s was timewise credible. We have proved that s is (pathwise) credible in Γ^* . ■

REMARK. It should be noted that not every assignment of dates to information sets will be valid even if the game is of perfect recall. For example, one cannot play the game of Figure 15 unless one is equipped with "time machine".

Without the dates this game makes perfect sense: Both players 2.1 and 3.1 have to act without knowing if the other one has already made his choice.

³¹ Because preferring to comply with s'' , when reached, means, preferring s'' to s . See also the proof of Lemma 7.2.

Figure 15. A game that is impossible to play



8. Other variants

In Section 3 we took the position that when an agent in the deviating coalition is indifferent to complying with s' or with s , he will not switch. This position gave rise to requirement (ii) in Definition 3.1. We could take the other position, that he will comply with s' in case of indifference. After all, what one promises to oneself is perhaps more important than what one promises to others... With this position in mind, an interesting variant may be to omit requirement (ii) altogether. Accordingly, we shall call a deviation s'_Q satisfying (i) and (iii) of Definition 3.1 an *optimistic credible deviation* and we define an *optimistic credible equilibrium* as an n -tuple of strategies from which there do not exist optimistic credible deviations.




Figure 16. An OCrE which is not closed

The set of optimistic credible equilibria (OCrE) coincides with NE in ordinary games, and the proof is essentially the same as in Theorem 4.1. Thus, OCrE is also an extension of the Nash solution concept. It, too, is not necessarily a closed set as the example in Figure 16 shows.

In this example $(M_1, L_2, [(1 - \varepsilon)L_3, \varepsilon R_3])$ is in OCrE whenever $0 < \varepsilon \leq 1$ but if $\varepsilon = 0$, agent 1.1 can instruct agent 1.2 to move right, and this is an optimistic credible deviation.

We believe that OCrE is worth studying, because of its simpler definition as compared to CrE and because it yields interesting strategies. The set OCrE need not contain APE and the game in Figure 17 is an appropriate example.




Figure 17. An example in which $\text{OCrE} \cap \text{APE} = \emptyset$

Here, (M_1, L_2, L_3) is the unique perfect equilibrium point for the agent form game. It is *not* optimistically credible because 1.1 can instruct 1.2 to move right. The point (M_1, R_2, L_3) is both in CrE and in OCrE.

Unfortunately, OCrE may be empty, as can be seen in the game of Figure 18.

In this example 1.2 and 2.1 play matching pennies. Thus, if $s \in \text{OCrE}$ then $s_{1.2} = (1/2, 1/2)$ and $s_{2.1} = (1/2, 1/2)$. But then, 1.1 can strike an optimistic credible deviation by instructing 1.2 to move “left”. One may feel that this counter example is sufficient to discard OCrE. We do not think so. OCrE conceptually is simpler than CrE and if it does contain some points, these may have some advantage.

In another variant of the OCrE concept, one requires that in case of indifference between complying with s' and complying with s in $\Gamma_{i,j}^{s'}$, $i, j \in Q$, agent i, j will comply with s' only




Figure 18. A game with an empty OCrE

if, *when the instruction was given*, agent $i.j$ was promised some gain in s' as compared to s . It is as if $i.j_0$ "tells" $i.j$ in Q : "If I would not have struck the deviation and we all played according to the original point s , you would have received a certain amount, a . Now that I strike the deviation, you will get b , and b is greater than a . True, when it will be your turn to play you will get b even if you play $s_{i,j}$, but where will your gratitude to me—your own flesh and blood—be for the profits I threw upon you?" We feel that this concept should also be studied, but we must say that it has one drawback: It requires utilities of agent $i.j$ to be defined even if s is played, perhaps under outcomes in which $i.j$ is not even created.

9. Other applications

Up to now we restricted ourselves to one scenario: The players are individuals who play a game in extensive form and their utilities may change during the play. But the model that we constructed may be useful in other applications. In this section we shall discuss two of them.

A. VIOLATION OF VON NEUMANN-MORGENSTERN INDEPENDENCE AXIOM. The last decade, which has witnessed a proliferation of studies of generalizations of von Neumann-Morgenstern [1947] expected utility theory, saw a revival of interest in the changing-preference problem. Indeed, in many models, the violation of von Neumann-Morgenstern independence axiom is equivalent to dynamic inconsistency in decision makers' preferences in a multi-stage decision problem under risk (see Hammond [1988], Karni and Safra [1989a, b], Machina [1989]). Differently put, the violation of the axiom is (at least technically) equivalent to a change in the decision-maker's utility function (over lotteries). Thus, our concept of CrE may be applied to such models as well.

B. WHEN A PLAYER REPRESENTS A GROUP OF INDIVIDUALS. In many game models, a player is not an individual. It can be a state, a political party, an organization, etc. In such cases it almost always makes no sense to attribute to such a player a von Neumann-Morgenstern utility function. We do not refer here only to situations that fall under Condorcet paradox (Condorcet [1785]), but to a more fundamental fact that a non-individual player *may not have a utility function*. What is the utility function of a state? The utility of the farmers? The utility of the manufacturers? The white collar people? The plain citizen? The prime minister? It is more natural to regard a player in such cases as an instruction-giving unit that *represents* groups³²—each endowed with its own utility function. Each such group is capable of deviating from the instructions given by the players and will deviate if it finds the deviation profitable and safe.

Applications of game theory in which players in the game model are not individuals are

³²Strictly speaking—individuals who sometimes are aggregated into groups to render further analysis feasible.

numerous, and among the most important. If the model developed in this paper could be used to handle them, that would be a significant contribution to applied game theory. For example, a game theorist could tell the instruction-giving unit: “You can, of course, decide on any strategy combination, but only credible combinations will be abided by. Here is the set of those combinations that are likely to be followed.” Unfortunately, the present model can handle only cases where perfect recall prevails.

For applications discussed here, perfect recall is not a reasonable requirement. Why should the farmers know what the manufacturers did in the past, for example?³³ Of course, if perfect recall does not prevail, one cannot limit oneself to behavioral strategies, but this is not the main source of difficulty. Take, for example, the one-person game of Figure 19.




Figure 19. A credible deviation that does not make sense

Here, (L_1, L_2) is not credible because agent 1.1 can instruct himself and agent 1.2 to move right. If agent 1.2 knew that agent 1.1 was moving right, he would have certainly chosen to comply. But he does not know, and if he complies, then it behooves agent 1.1 to remain at L_1 so as to get 3. So, the deviation (R_1, R_2) is perhaps not safe, in spite of the fact that it complies with Definition 3.1.

³³Of course, if the instructions can only be delivered publicly, as it should be in most cases of democratic countries, we can claim that perfect recall prevails.

If the game lacks perfect recall, difficulties start with the construction of the model itself. The model should allow an agent to occupy several information sets, belonging to different dates, if dates are specified. It should be clear who plays after whom, so as to know to whom instructions of the deviating player can be given. The strategy space should also be clearly specified: Do we allow agents to use correlated strategies, for example? Suppose the agents have agreed on a strategy combination and that some agents deviated in such a way that others do not know it, due to the imperfect recall. And suppose that, contrary to the agreement, an agent finds himself at an information set that should have been reached with zero probability. He knows that a violation occurred but often does not know which one. How will he interpret the observation? We see that with lack of perfect recall all the problematic issues of “refinements” pop up in spite of the fact that our goal was to only generalize the Nash solution. Thus, applications of CE to games where a player represents several groups is at present quite limited. The extension of this model to games without perfect recall remains a challenging topic for further research.

REFERENCES

- Allais, M., "Economie et Interet," Imprimerie Nationale, Paris, 1947.
- Aumann, R. J., *What is Game Theory Trying to Accomplish*, in "Frontiers of Economics," K. J. Arrow and S. Honkapohja (eds.), Basil Blackwell, Oxford, UK-New York, 1985, pp. 28-87.
- Aumann, R. J., and A. Brandenburger, *Epistemic Conditions for Nash Equilibrium*, RM 91-042 Harvard Business School, 1991.
- Bernheim, B.D., B. Peleg, and M. D. Whinston, *Coalition-Proof Nash Equilibria, I. Concepts*, Journal of Economic Theory **42** (1987), 1-12.
- Blackorby, C., D. Nissen, D. Primont, and R. R. Russell, *Consistent Intertemporal Decision-Making*, Review of Economic Studies **40** (1973), 239-248.
- Condorcet, M.J.A. (de Caritat, Marquis de), "Essai sur l'application de l'analyse à la probabilité de décisions rendues à la pluralité de voix," Paris, 1785.
- Hammond, P. J., *Changing Tastes and Coherent Dynamic Choice*, Review of Economic Studies **43** (1976), 159-173.
- _____, *Consequentialist Foundations for Expected Utility*, Theory and Decision **25** (1988), 25-78.
- Harsanyi, J. C., *Welfare Economics of Variable Tastes*, Review of Economic Studies **21** (1953), 204-213.
- _____, *Games with Incomplete Information Played by "Bayesian" Players*, Parts I, II, III, Management Science **14** (1968), 159-182, 320-332, 468-502.
- Karni, E., and Z. Safra, *Dynamic Consistency, Revelations in Auctions and the Structure of Preferences*, Review of Economic Studies **56** (1989a), 421-434.
- _____, *Ascending Bid Auctions with Behaviorally Consistent Bidders*, Annals of Operations Research **19** (1989b), 435-446.
- Kuhn, H. W., *Extensive Games and the Problem of Information*, in "Contribution to the Theory of Games," Annals of Mathematics Studies Vol II, H. W. Kuhn and A. W. Tucker (eds), 1953, pp. 193-216.
- Machina, M., *Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty*, Journal of Economic Theory **27** (1989), 1622-1668.
- Nash, J. F., Jr., *Non-cooperative Games*, Annals of Mathematics **54** (1951), 289-295.
- Peleg, B., *On Perfectly Coalition-Proof Nash Equilibria*, in "Equilibrium and Dynamics: Essays in Honor of David Gale," W. Brock and M. Majumdar, eds. (to appear).
- Peleg, B., and M. E. Yaari, *On the Existence of a Consistent Course of Action when Tastes are Changing*, Review of Economic Studies **40** (1973), 391-401.
- Phelps, E. S., and R. A. Pollak, *On Second-Best National Saving and Game-Equilibrium Growth*, Review of Economic Studies **35** (1968), 185-200.
- Pollak, R. A., *Consistent Planning*, Review of Economic Studies **35** (1968), 201-208.
- _____, *Habit Formation and Dynamic Demand Functions*, Journal of Political Economics **78** (1970), 745-764.
- _____, *Habit Formation and Long-Run Utility Functions*, Journal of Economic Theory **13** (1976), 272-297.
- Savage, L. J., "The Foundations of Statistics," Second Revised Edition, Dover Publications, New York, 1972.
- Schoeffler, S., *Note on Modern Welfare Economics*, American Economic Review (1952), 880-887.
- Selten, R., *Re-examination of the Perfectness Concept for Equilibrium Points in Extensive Games*, International Journal of Game Theory **4** (1975), 25-55.
- Strotz, R. H., *Myopia and Inconsistency in Dynamic Utility Maximization*, Review of Economic Studies **23** (1956), 165-180.
- Von Neumann, J., and O. Morgenstern, "Theory of Games and Economic Behavior," Princeton University Press, Princeton, New Jersey, 1947.
- Von Weizsaecker, C. C., *Notes on Endogenous Changes of Tastes*, Journal of Economic Theory **4** (1971), 345-372.

æ

Keywords. games, non-cooperative games, extensive form, utility theory, Nash equilibrium, changing preferences, changing utilities.