Formation of Preferences and Strategic Analysis: Can they be De-coupled?

Nimrod Megiddo

September 24, 2003

1 Introduction

State-of-the-art computers and network technologies have already enabled the use of intelligent software agents for making decisions in real time. Software agents interact both with humans and with other such software agents. Many types of online interaction seem to be amenable to game-theoretic analysis, since they involve both competition and cooperation. For example, agents bargain with humans and with other agents about prices, participate in auctions, coordinate actions and utilization of resources, and also establish and manage trust. The question often arises whether game theory, as we know it, can be successfully applied in such situations. Theoretically, there are numerous ways in which game theory might be applicable. In this paper, we are interested only in the type of application, where players are programmed to behave in specific situations, based on game-theoretic analysis of preferences and choice of actions. The amounts of resources consumed by a computer program before it makes a decision must be limited by the operational requirements of the system. Therefore, the analysis must be relatively simple. Also, there has be a way to map such analysis to a computational procedure, which is essentially sequential. So, the high-level question we are interested in is whether there is a way to endow computer programs with game-theoretic rationality.

In principle, game theory could be applied in at least two kinds of circumstances. The first is the single game situation: a software engineer may wish to develop an "agent" for playing a *particular* game. In this case, the agent amounts to an implementation of a single strategy for playing the particular game. The software engineer would have to study the particular situation, take into account the preferences of all the participants and then design a good strategy for the agent which, in turn, would be reduced to a computer program. Obviously, the software engineer would have to be familiar with game theory in order to apply it in the particular situation under consideration. Most software engineers, however, simply apply common sense in such situations. The second kind of application is the more challenging problem of applying game theory as a *methodology*. Imagine an agent whose task is to play various games. For the sake of concreteness, suppose an agent has to play games in extensive form. Thus, the agent would receive a description of a game as a tree with information sets associated with players, and outcomes specified at the leaves. Here, we do *not* assume that the utility values, which various players attach to the outcomes, are known. As part of the task of playing the game, the agent would at least have to form its own preferences over the set of outcomes, let alone beliefs about the preferences of other players.

Even though we have mentioned software agents repeatedly, their role in our discussion is not essential. We allude to a software agent because in this way we can highlight the need for a precise prescription for playing a game. Obviously, computer programs perform their computation one step at a time. Thus, in order to specify how a computer program should play various games, a formal and precise prescription has to be furnished as to what steps are taken at what times. A prescription is a step-by-step procedure that works for any game specified in a certain format. It has to analyze the game and develop a strategy of how to play it. This procedure has to be practical in the sense that it should run in reasonable time and take reasonable amount of space.

Any game-theoretic analysis must rely on the preferences of players. Since in practice the preferences are *not* given together with the game, they have to be elicited from players if possible, or derived according to some methodology. The outcomes of the game may, of course, give some hints about the preference order. For example, monetary value, in general, is correlated with utility. Of course, in some cases the description of the game may include the exact preferences. But in this paper we are interested only in games where the exact preferences have yet to be determined. The question here is whether it is a valid approach to fix the preferences first, and only later analyze the resulting game with fixed preferences, or rather analyze preferences and strategies simultaneously; if so, how should this be done? It seems that the traditional approach in game theory and decision theory is that elicitation of preferences should precede the choice of actions. We demonstrate in this paper that, in certain games, the formation of preferences over outcomes and the strategic analysis of the game are intertwined to the extent that common resolution procedures, such as backward and forward induction, cannot be applied.

The question raised here pertains to the applicability of game theory to concrete real-life situations. Our interest here is in the normative rather than the descriptive aspect of game theory, namely, how should a game be analyzed for play in the real world when the preferences have not been fully formed. The common approach of game theory to situations where some preferences are not known is to view the game as one of incomplete information [3] and expand it into much larger game with complete information. We discuss this approach below. Other seemingly related topics are psychological games [2] and games with utility changing during the play [1]. We discuss the connections to these topics later.

2 Changing preferences in repeated games

2.1 The traditional approach

Many papers have been written about repeated games. Intuitively, a repeated game means playing repeatedly the same single stage game with some monitoring of payoffs or actions. This obviously means that the possible *physical* outcomes are the same in each stage. But what about *utility payoffs?* Can they be the same in each stage? In standard game theory, the description of a (complete information) game includes the final utility payoffs. Therefore, the final utility payoff in a standard repeated game is usually defined as an aggregate of the utility payoffs obtained during the various stages of the game. In particular, the possible utility payoffs in a single stage of the repeated game are independent of the play in the previous stages; in other words, they are the same in every stage. In our opinion, this assumption is not realistic, as we attempt to demonstrate using the repeated games of Prisoner's Dilemma (PD) and the Battle of the Sexes (BS).

2.2 Repeated Prisoner's Dilemma

In a single stage of the PD game, each player can choose either to "cooperate" (C) or to "defect" (D). Consider, for instance, a single stage PD game, where the players are each paid \$100 if both cooperate, \$20 if both defect and, otherwise, the defector gets \$110 while the cooperator gets zero. Obviously, the preferences of the players should be defined over the set of pairs

$$\{(\$100, \$100),\ (\$20, \$20),\ (\$0, \$110),\ (\$110, \$0)\}$$

or, equivalently, the set of action pairs

$$\{(C, C), (D, D), (C, D), (D, C)\}\ ,$$

rather than the set of dollar amounts

$$\{\$100,\ \$20,\ \$0,\ \$110\}\ .$$

This distinction is required if a player cares not only about his own dollar payoffs but also about the other player's dollar payoffs. Obviously, the dilemma prevails[1] if the utilities are monotone in terms of these dollar amounts, i.e.,

$$u_1(D, C) > u_1(C, C) > u_1(D, D) > u_1(C, D)$$

and

$$u_2(C, D) > u_2(C, C) > u_1(D, D) > u_1(D, C)$$

where $u_i(\cdot, \cdot)$ is the utility function of player i ($i = 1, 2$). If it is accepted that a player may care about the other player's dollar payoffs, then it is hard to justify an assumption that this attitude is *independent* of the play in previous stages of the game.

Consider, for example, the PD game being played twice with the same dollar payoffs as above. After the first round, the choices of both players

[1]Defection is a dominant strategy for each player, but it is better for both players if both cooperate than if both defect.

are disclosed, and then the second round is played. Obviously, after the first round, the attitudes of the players toward each other may change in view of the their respective choices in the first round. For example, if the play in the first round were (C, D) then player 1 might be angry at player 2, whereas he might feel remorseful if the play in the first round were (D, C). These different feelings may affect the preference order of player 1 in the second round. Suppose the utility payoffs in the first round are: $u_1^1(D, C) = u_2^1(C, D) = 105$, $u_1^1(C, C) = u_2^1(C, C) = 100$, $u_1^1(D, D) = u_2^1(D, D) = 20$ and $u_1^1(C, D) = u_2^1(D, C) = -10$. The utility payoffs in the second round may have four different sets of values, corresponding to the four different possible ways the first round could be played. For $x, y, v, w \in \{C, D\}$ and $i = 1, 2$, denote by $u_i^2(x, y; v, w)$ the utility payoff of player i in the second round if the play in the second round is (x, y), given that the play in the first round was (v, w). Suppose the play in the second stage is (D, C). It would be quite natural for player 1 to have $u_1^2(D, C; D, C) < u_1^2(D, C; C, D)$, since the play $(D, C; C, D)$ could be interpreted as if player 1 rightfully "punishes" player 2 for defecting in the first stage, whereas in $(D, C; D, C)$ player 1 "betrays" player 2 who was trying to be cooperative. Suppose, for example,

$$u_1^2(D, C; C, C) = u_2^2(C, D; C, C) < u_1^2(C, C; C, C) = u_2^2(C, C; C, C) = u_1^1(C, C) \,,$$

and in all other cases, $u_i^2(x, y; v, w) = u_i^1(x, y)$. It follows that, subsequent to a play of (C, C) in the first round, the cooperative play in the second round is an equilibrium. Therefore, it is an equilibrium in the two-round game for both players to play Tit-for-Tat[2], resulting in a cooperative play in both rounds.

In order to decide what to do in the first round, a rational player would have to form preferences over the set of possible outcomes $\mathcal{O} = \{(C, C)$, (D, D),(C, D),$(D, C)\}$ in the second round, given any possible play in the first round. The player has to consider playing the second round in four different contexts rather than merely form preferences over the set $\mathcal{O}$. Thus, he may come up with preferences such as $u_1^2(D, C; C, C) < u_1^2(C, C; C, C)$, whereas $u_1^2(D, C; C, D) > u_1^2(C, C; C, D)$.

Suppose in a certain context in the second round the utility payoffs are no longer monotone in terms of the dollar payoffs. Then, defection in that

[2]The Tit-for-Tat strategy is to play C in the first round, and in the second round to play whatever the other player has played in the first round.

particular context may no longer be a dominant strategy. It is plausible that an equilibrium in that context would require mixed strategies. Another possibility is that, when the equilibrium payoffs from the second round are added to the payoff matrix of the first round, then defection in the first round is no longer a dominant strategy, and an equilibrium in the two-round game would require randomization even in the first round.

The two-round PD is not too complicated, since the second round may be played only in four different contexts. Obviously, in general, the $(k+1)^{st}$ round of the repeated PD may be played in 4^k different contexts. Thus, at least in principle, a rational player would have to form preferences over the set $\mathcal{O}$ in 4^k different contexts. Moreover, since the final utility payoff is an aggregate of the utility payoffs of the various rounds, the player would have to form cardinal, rather than just ordinal, utility functions for the various contexts, so that these utility payoffs could be aggregated meaningfully over the rounds of the game. Of course, in order to analyze the game properly, each player would also have to form in his mind such utilities for the other player. Unlike more problematic examples discussed in this paper, it does not seem that a player in the two-round PD would have too much difficulty in forming an opinion about the other player after the first round, since the second round is not too complicated. As the more problematic examples show, a player may face more difficulty in interpreting the other player's actions when he does not know where to start the analysis.

The difficulties discussed above raise the question whether any repeated game with certain physical outcomes could at all be analyzed as a repeated game with the same utility payoffs in each round. It is conceivable that a player in a repeated PD would be able to form reasonable preferences as a function of the play in the previous rounds. Instead of considering all the possible 4^k different contexts explicitly, the player's preferences may depend only on some "features" of the context rather than its full detail. In other words, the player might implicitly consider a small number of equivalence classes of contexts. We explain this issue in the next example, using the game of Battle of the Sexes (BS).

2.3 Repeated Battle of the Sexes

In a single round of the BS game, each of Husband and Wife has to choose either H or W, i.e., going to Husband's favorite place or Wife's favorite place.

Imagine this game being simulated in a laboratory, where the subjects receive the following dollar payoffs: $P_H(H, H) = P_W(W, W) =$ \$200, $P_H(W, W) = P_W(H, H) =$ \$100, $P_H(H, W) = P_W(H, W) =$ \$10 and $P_H(W, H) = P_W(W, H)$ = \$0. These dollar payoffs are supposed to reflect that both players prefer to go to the same place rather than go to different places, but each has a different favorite place, which he or she prefers that both would go to.[3] Consider a two-round BS game and consider a situation where in the first round the players chose (H, H). Reasonable players might expect by reason of "fairness" the play in the second round to be (W, W). For $x, y, v, w \in \{H, W\}$ and $i \in \{H, W\}$, let us denote by $u_i^2(x, y; v, w)$ the utility payoff of player i in the second round, given that the play in the first round was (v, w) and the play in the second round is (x, y). It would be quite natural for Wife to have $u_W^2(H, W; H, H) < u_W^2(H, W; H, W)$, since subsequent to (H, H) she expects Husband to play W, whereas subsequent to (H, W) her expectations are not clear; in the latter case, a play of (H, W) could be more acceptable than in the former one. Also, it would be quite natural to have $u_W^2(H, W; W, H) > u_W^2(H, W; W, W)$ since in the case (W, H) was played in the first round, Wife may feel she has the "right" to play W in the second round, whereas in the case (W, W) was played in the first round she may feel it would not be nice of her to play W again in the second round.

3 Dependence of utility on actions not taken

In a typical tree representation of a game, the same "physical" outcome may be associated with multiple leaves. In particular, such redundancy occurs when the game is defined by a concise set of rules and the tree enumerates all the possible plays. A naïve approach is to determine preferences over the set of distinct physical outcomes, but it is well known that this approach certainly does not suffice for capturing what really happens in a game. In principle, each leaf of the game tree represents a different outcome and hence the preferences should be determined over the set of leaves. The latter viewpoint may be called the "common" approach. Of course, it can be used in practice only when the size of the tree is manageable. However, we wish to address here a different issue. The question we are concerned with is whether

[3] In the literature, the utilities are usually assumed to satisfy $u_H(H, W) = u_H(W, H)$ and $u_W(W, H) = u_W(H, W)$, but this is only a minor difference.

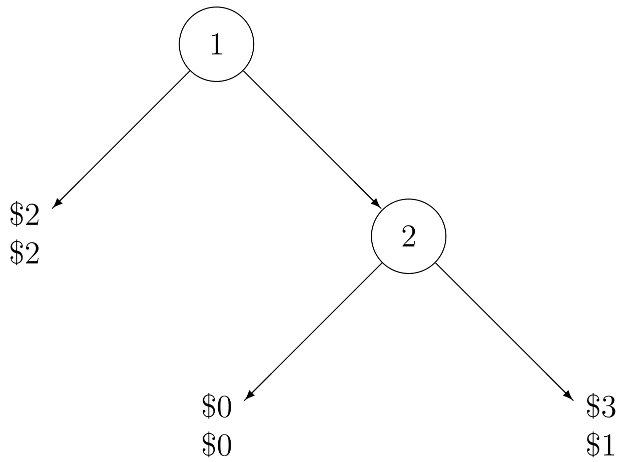

Figure 1: A simplified ultimatum game

the step of forming such preferences over the leaves of the tree can precede the analysis of the game. In this section we present a few examples, which *can* be handled by the common approach. These examples are presented merely for the sake of highlighting the underlying issues. Later, we explain how these examples can be extended into more problematic situations.

In the game depicted in Fig. 1, player 1 moves first, choosing between (i) the outcome of \$2 payoff to each player, and (ii) letting player 2 choose. If player 2 is called upon to play, he has to choose between zero payoff to each player and the payoff of \$3 to player 1 and \$1 to himself. Thus, the distinct physical outcomes in this game are: (\$2, \$2), (\$0, \$0) and (\$3, \$1). Consider another game, where player 1 chooses between (i) zero payoff to each player, and (ii) letting player 2 choose between (\$2, \$2) and (\$3, \$1). Thus the physical outcomes in both games are the same. In the former game, it is plausible that player 2 would prefer (\$0, \$0) over (\$3, \$1), but there does not seem to be any reason for that to happen in the latter one. Returning to the former game, the preference of player 2 between (\$0, \$0) and (\$3, \$1) seems to depend on the fact that if indeed he has to choose between the two, then it is because player 1 has rejected the (\$2, \$2) outcome. That choice does affect how player 2 evaluates the benefit to player 1. He might be upset to the extent that he would choose (\$0, \$0) rather than (\$3, \$1) in order to punish player 1. But if instead of (\$2, \$2) the rejected outcome were, say,

(\$1, \$1), then there would not seem to be any reason for player 2 to prefer the (\$0, \$0) over (\$3, \$1). So, the outcomes that could have been reached but have already been rejected, affect the preferences over outcomes that could still be reached.

The lack of knowledge of the references of other players is traditionally treated by distinguishing "types" of players. Let us first recall this approach. In traditional game theory, the analysis of a game relies is based upon individual utility payoffs associated with the various outcomes of the game. In an extensive form game, each player is assumed to have a "utility function" which maps leaves of the game tree to real numbers. The game is said to be one of complete information if the utility payoffs of all players are common knowledge among all players; otherwise, it is a game of incomplete information. Harsanyi [3] proposed a remedy for the lack of common knowledge of the utility payoffs. He proposed to expand the game so as to account, probabilistically, for the various possible types of players. The types are jointly drawn from some commonly known prior probability distribution, and each player is privately informed of his own type. It is not clear, in general, how the common prior distribution arises. In practice, it seems difficult to fix a suitable set of types without first analyzing the game strategically. Harsanyi's proposal replaces a game of incomplete information by one of complete information, but this replacement requires first fixing the set of types and the common prior joint probability distribution over them.

Returning to our example above, within a Bayesian framework, player 1 would have a certain probability p that player 2 would prefer (\$0, \$0) over (\$3, \$1), given that player 1 has rejected the (\$2, \$2) outcome (see Fig. 2). Denote by $u_1(\cdot,\cdot)$ player 1's utility function. Thus, player 1 should reject the (\$2, \$2) outcome if and only if $u_1(\$2,\ \$2) < p\,u_1(\$0,\ \$0) + (1-p)\,u_1(\$3,\ \$1)$. However, the utility of player 1 from each of the outcomes must reflect the *context* in which the outcome results rather than its mere monetary value. For example, if the outcome (\$3, \$1) is reached, it must also involve a bit of anger on the part of player 2, which would result in some utility value for player 1, depending on player 1's own personality type.

Perhaps a more convincing example is depicted in Fig. 3. Consider the question of player 1's preference between the outcomes (\$2, \$101) and (\$1, \$0). The preference is relevant if and only if player 1 is called upon to choose between the two. The context in which this event occurs is as follows. Player 1 has previously chosen not to terminate the game with the outcome of (\$10,

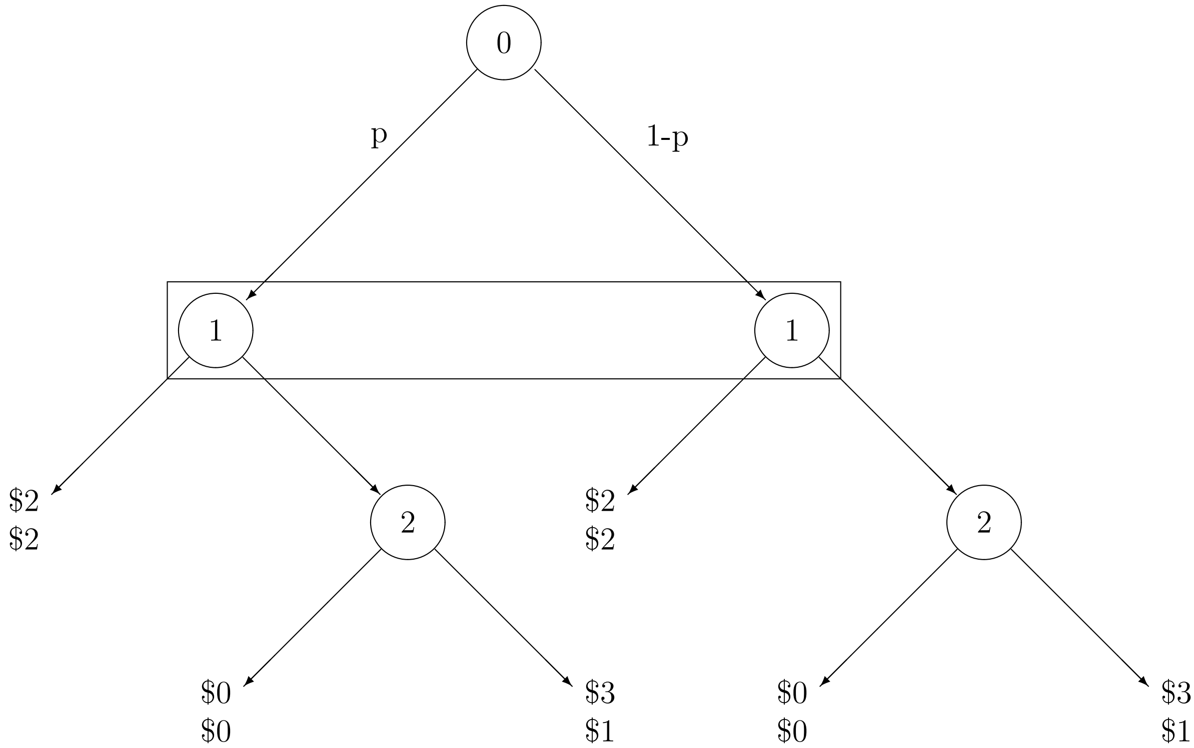

Figure 2: A Bayesian expansion of the simplified ultimatum game

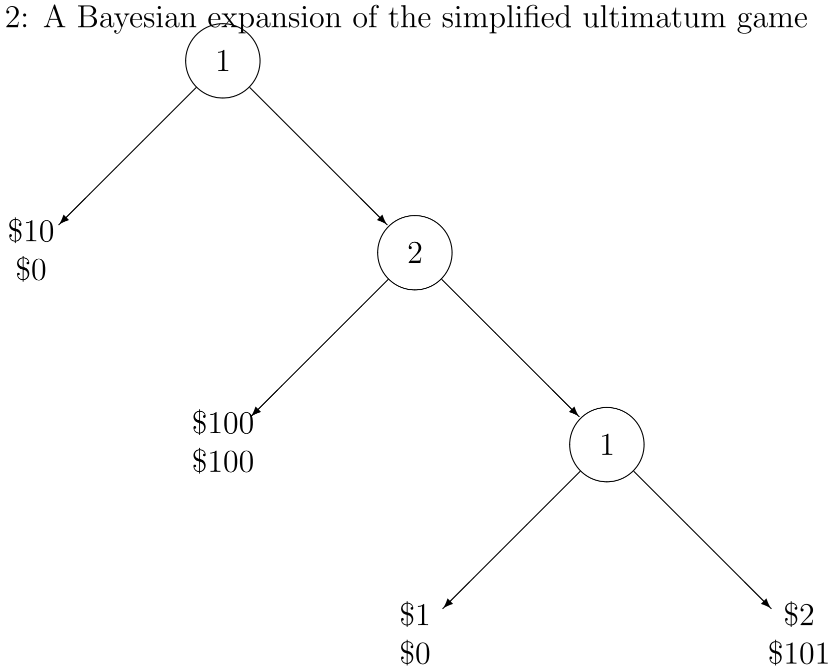

Figure 3: Punishing a greedy player

\$0) but rather let player 2 make a choice. As a consequence, player 2 had the option to terminate the game with the outcome of (\$100, \$100) but has rather chosen to let player 1 choose between (\$2, \$101) and (\$1, \$0). This action of player 2's might be seen by player 1 as quite selfish for reasons as follows. First, it was player 1 who made it possible for player 2 to get a non-zero payoff, and player 2 could have secured a payoff \$100 for each of the players. But, instead, it seems that player 2 has aimed at getting an extra single dollar for himself while player 1 getting \$2, assuming that player 1 would eventually prefer \$2 over \$1, regardless of what player 2 gets. Moreover, player 2 knew that player 1 could have easily gotten \$10, so his only reason for not terminating the game with that outcome must have been the hope for the payoff \$100 to each. So, it would not be too surprising if at this point player 1 indeed preferred (\$1, \$0) over (\$2, \$101). Note that this preference arises in a specific context of actions by both players. This preference is a result of a certain interpretation of actions, which causes player 1 to evaluate player 2's personality in a certain way. The argument here does not rely on beliefs of player 1 about player 2, because the revelation has already occurred. It is simply the unpleasant personality of player 2, revealed by the latter's actions, which causes player 1 to sacrifice some payoff in order to teach player 2 a "lesson." If, instead of (\$10, \$0), the payoff were, say, (\$10, \$200), then the argument could yet be quite different.

Another example of the effect of a "revealed personality" is depicted in Fig. 4. Here, there are three players but player 3 never makes any decision. Player 2 has to choose between (\$100, \$10, \$0) and (\$200, \$9, \$0), after player 1 has chosen between (\$10, \$0, \$100) and (\$9, \$0, \$200). Player 2 might be "nice" and would prefer (\$200, \$9, \$0) over (\$100, \$10, \$0) without any prior information about player 1's personality. If, however, player 1 has chosen (\$10, \$0, \$100) rather than (\$9, \$0, \$200), then player 1 himself does not seem to have acted nicely toward player 3, so player 2 might not be nice to him.

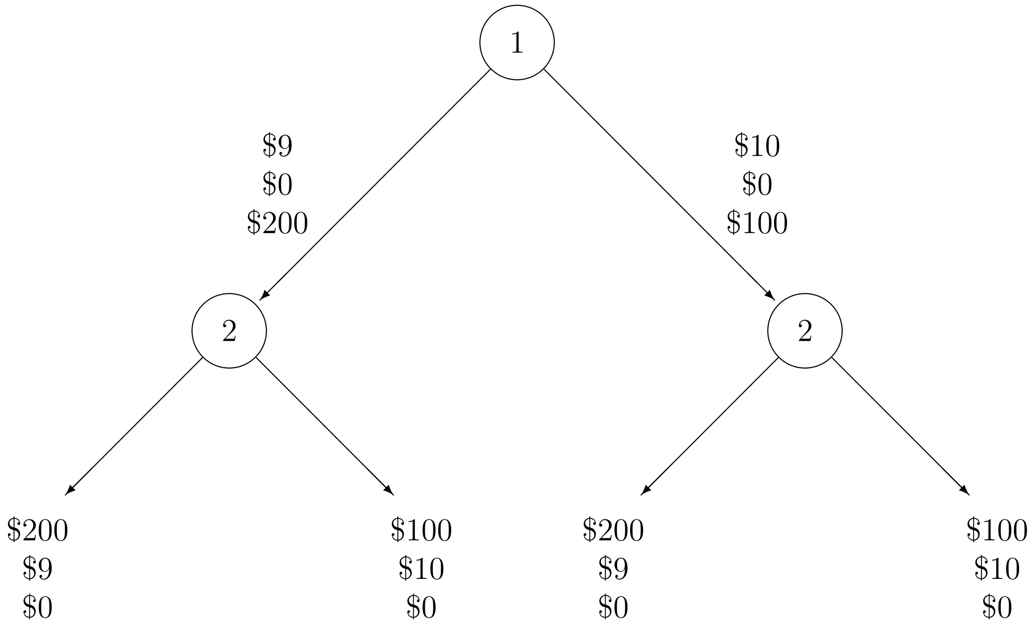

Figure 4: Learning from a player's treatment of a third party

4 The problem of simultaneous valuation of sub-games

The main reason why the examples of the preceding section *could* be handled by the traditional game theory approach is that *they do not require simultaneous valuation of several sub-games.* For example, in the game depicted in Fig. 3, player 2 can terminate the play with definite payoffs of ($100, $100). A more complicated situation would arise if the same move lead to a sub-game for which payoffs of ($100, $100) could only be anticipated by some analysis rather than guaranteed as the final result. In the game of Fig. 3, it is relatively easy to interpret a choice of player 2 not to go left, as a manifestation of greed and lack of gratitude to player 1's, whose previous choice to go right seems to benefit player 2. Thus, the analysis amounts to understanding the beliefs of both players with regard to the type of player 2, possibly including beliefs about beliefs, and so on. In particular, it is not too complicated for player 1 to formulate his preferences at the second decision node in view of player 2's choice to go right. On the other hand, if the node ($100, $100) were replaced by a certain sub-game, then in order for player 1 to be able to interpret player 2's decision to avoid that sub-game, player 1 would have to *analyze* that sub-game in order to interpret player 2's choice. Of course, player 2 would also need to carry out such analysis in order to reason about player 1's anticipated reaction. For example, suppose the sub-game is the one shown in Fig. 5. This game is the same as the sub-game of the game of Fig. 3, after player 1 moves right, except that where the roles of players 1 and 2 are reversed. In order to valuate the sub-game of Fig. 5, one would have to develop an interpretation of a possible a decision by player 2 to enter this sub-game (by moving left in the game of Fig. 3) rather than let player 1 choose between ($1, $0) and ($2, $101) as the final outcome.

At this point it should be quite clear that, in general, a sub-game has to be evaluated within the context in which it is reached. The context is important because it may reveal information about the personalities of players and therefore affect the *utility values* associated with outcomes of the sub-game. In particular, the valuation of opportunities, which are missed by choosing to play one sub-game rather than another one, affects the valuation of the sub-game in question. This fact gives rise to a *circularity* in the analysis since the valuation of one sub-game affects the valuation of the a second sub-

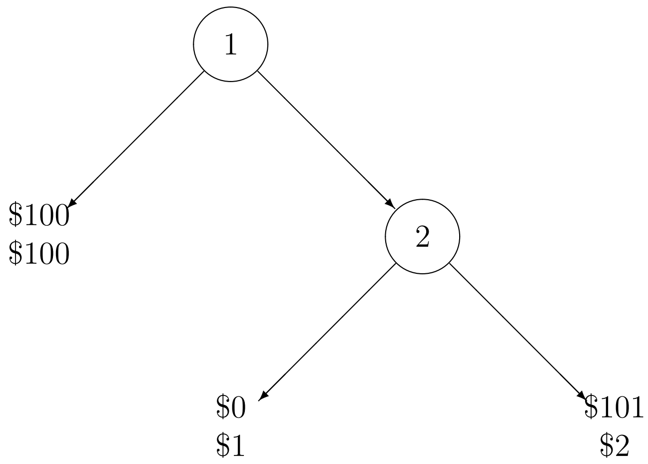

Figure 5: A sub-game

game and vice versa. This is one reason why the formation and preferences and strategies analysis are difficult to de-couple. To illustrate this problem, consider a situation where in the first move player 1 chooses which sub-game, L or R, is about to be played by him and player 2 (see Fig. 6), and suppose these sub-games are non-trivial. Thus, player 1 has to form an opinion as to which sub-game he prefers to play. We argue that *player 1 cannot arrive at such preference by first analyzing one of the sub-games and then the other one.* Suppose player 1 attempts to analyze the sub-game L first. One important issue, which arises in the sub-game L, is the anticipated reaction of player 2 to player 1's choosing not to play the sub-game R. In order to interpret such a choice properly, player 2 (and therefore also player 1) would have to analyze the sub-game R in order to see, for example, whether player 1 has acted generously or selfishly. This distinction is crucial because it affects the way player 2 may be willing to play the sub-game L. Of course, this requirement to analyze the sub-game R contradicts the assumption that the sub-game L can be analyzed first. It turns out that, unless one of the sub-games is trivial, the valuation of the two sub-games has to be carried out *simultaneously.* In Section 6.2 we describe a concrete example, where the alleged circularity is demonstrated.

The discussion above suggests that, in general, games cannot be analyzed by considering sub-games sequentially and, at best, one perhaps can only

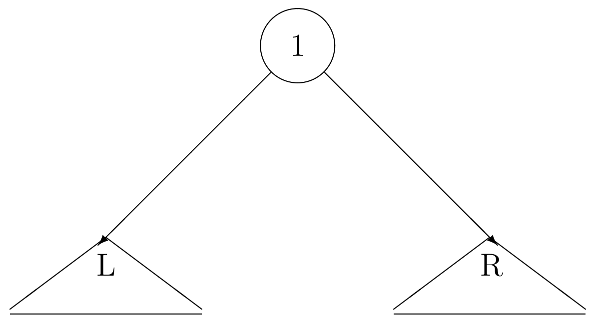

Figure 6: A symmetric circularity

achieve some kind of consistency between preferences and choices of strategies, while preferences may not be unique. In the next section we propose one possible notion of such consistency.

5 Consistency of preferences and choice of action

In traditional game theory, each player is assumed to have a unique preference order over the set of probability distributions over outcomes of the game. The reasoning about behavior in the game relies on this assumption, so players need to form their preferences before the analysis of the game begins. Our point of departure from traditional game theory is related to this assumption of uniqueness of preferences. We question the assumption that, prior to the analysis of the game, players must have unique preference orders over the outcomes of the game. In particular, in our view, the existence of a unique von Neumann - Morgenstern utility function $u_i : \mathcal{O} \rightarrow \Re$, which necessitates the existence of unique preferences over probability distributions over outcomes, cannot always be assumed a priori. On the other hand, the *analysis* of the game may eventually lead to a unique utility function. We argue that the choices of strategies and preferences though have to be *consistent* in a sense which we introduce below.

The puzzling question is as follows. Assuming the preferences have not been fixed yet, how does a player start reasoning about the game. In most cases, a player would probably start with a proper subset of the possible preference orders or even utility functions over the set $\mathcal{O}$ of outcomes. Thus, the game could, at least in principle, be analyzed with respect to every plausible tuple of utility functions. Furthermore, a player may have prior beliefs about which preference order (or utility function) would result from a careful analysis of the game. The set of a priori possible utility functions of a player may reflect the player's character or moral values. Let us denote by $U^i \subset \Re^{\mathcal{O}}$ the set of a priori plausible utility functions of player i. First, without loss of generality, we may assume that for each $u \in U^i$, $0 \leq u_j \leq 1$ for every $j \in \mathcal{O}$, so U^i is compact. Second, we may assume that U^i is convex. We do not attempt to justify this assumption here. We leave it to the reader to decide whether this assumption is justified. One possibility is that U^i is the set of all utility functions which are consistent with a certain partial preference order. Thus, the player may have a partial order but still not be sure about the exact preference order between every two alternatives. Another possibility is that the player may have a set of conditions which the utility function must satisfy but these conditions do not define a unique function.

In view of our discussion so far, it seems that a "solution" of a game should include not only choices of strategies but also choices of utility functions of the players. It is natural to generalize the Nash-equilibrium concept to the situation described here, namely, a solution of the game would consist of particular utility functions $u^i \in U^i$ and strategies σ^i, which are in equilibrium with respect to the fixed u^is. But without further requirements, there is no restriction on the chosen u^is. Thus, we also require a certain consistency of the chosen utility functions with respect to the actions of the players. More specifically, suppose each player has a "utility revision map" Ψ_i reflecting the player's thought process as follows. Denote by $U = U^1 \times \cdots \times U^n$ the cartesian product of the sets of plausible utility functions, and denote by $\Sigma = \Sigma^1 \times \cdots \times \Sigma^n$ the cartesian product of the strategy spaces. While attempting to form preferences and decide what to do, a player may consider some pair $(u, \sigma) \in U \times \Sigma$ and try to decide whether such a pair is a reasonable outcome of the game. Our examples show that a player may revise his preferences over outcomes in view of the behavior of other players, since his attitude to each of the other players may change. Thus, given a pair (u, σ), a player i may change his utility function from u^i to another function v^i. This observation gives rise to the concept of a *utility revision map* (URM). We assume that, in principle, each player i has a map $\Psi_i : U \times \Sigma \to U^i$, which gives the revised utility function $v_i = \Psi_i(u, \sigma)$ of player i in case the choices of strategies and utility functions are (u, σ). When the URMs are combined, we can talk about the utility revision map $\Psi : U \times \Sigma \to U$ which models the re-evaluation of outcomes that players may apply in view of the choices other players make and their previously estimated utilities. Thus, a utility value, which is assigned tentatively to an outcome, may be revised if the choices of strategies are revealed, because the latter may change the attitudes of players towards one another, and hence the utilities they attach to outcomes affecting other players. The reason why the revised utilities may depend on the current utilities is that one player's interpretation of the actions of another player depends on the how the other player seems to value various outcomes. Thus, even though the strategies of the players may be in equilibrium, some player may still think he is not playing a best response because his utility function has changed.

We say that a utility profile $u \in U$ is a *relative fixed-point* with respect to a strategy profile $\sigma \in \Sigma$ if $u = \Psi(u, \sigma)$. Thus, if the players chose the strategy profile σ and assumed the utility profile to be u, then they would not revise this utility profile. It is well known in game theory that for any fixed utility profile u, there exists a continuous map $\Phi_u : \Sigma \to \Sigma$, which can be used for proving the existence of a Nash-equilibrium (see, for example, [4]). Thus, for in order to formalize consistence of preferences and actions we may consider the following map $\Theta : U \times \Sigma \to U \times \Sigma$, defined by $\Theta(u, \sigma) = (u', \sigma')$ where $u' = \Psi(u, \sigma)$ and $\sigma' = \Phi_u(\sigma)$. The existence of a fixed point of Θ under our assumptions is obvious by Brouwer's theorem. Such a fixed point is a consistent combination of plausible preferences and strategies.

6 On types of players

6.1 A "simple" example

Consider the following two-person game where the players are given two desirable objects, one of which is a bit more desirable than the other, and they have to decide who gets which object. Let us refer to these objects as the "large" and the "small". The assumption is that, if the choice had to be made outside the current game, then each player would rather have the large object while the other player is getting the small one. Player 1 moves first and has only two options: (i) terminate the game by letting a referee toss a fair coin to determine who gets the large and who gets the small, and (ii) let player 2 make choose of who gets which object. If player 2 is called upon to play, then his only choice is between taking the large object for himself (and leaving the small one to player 1), or taking the small one for himself (and leaving the large one for player 1). There could be a number of reasons why player 1 may choose to let player 2 make the decision. Here are some of them:

1. Player 1 might be "nice" and would enjoy letting player 2 have the larger, more desirable object.

2. Player 1 might not be so nice and would prefer getting the large object while player 2 gets the smaller one, but he believes that player 2 is nice and would let player 1 have the large object.

3. Player 1 might prefer to let player 2 "participate" in the game, and this might be more significant than the actual object received.

If player 2 were called upon to play, then his decision would depend both on his basic preferences and on his beliefs about player 1. Thus, there could be several possibilities:

1. Player 2 might think that player 1 was really nice, and therefore player 2 would react nicely and let player 1 have the large object.

2. Player 2 might suspect that player 1 actually expected player 2 to grant him the large object, and therefore player 2 would not be willing to do that.

3. Player 2's decision might be independent of what he might believe about player 1, i.e., player 2 might be nice or not and that property alone would determine his choice.

Clearly, there are more than two "outcomes" in this game even though there are only two possible assignments of the two objects to the two players. A simple tree describing this game has four leaves, but if each player can be one of two types, then a larger tree with sixteen leaves would be required.

6.2 Modelling with "types" requires strategic analysis

Consider one more example. Suppose player 1 picks one of two sub-games (see Fig. 7). In the Left sub-game, player 2 has to choose between ($0, $49) and ($1100, $50). So, if 2 prefers ($0, $49) over ($1100, $50), then it might be interpreted as if he wishes to "punish" player 1 for something. So, in order for player 2 to clarify for himself whether he wishes at all to punish player 1, player 2 has to consider the Right sub-game and imagine why player 1 has chosen not to play Right. In the Right sub-game player 2 has to choose between ($1000, $90) and ($0, $100). If he prefers ($1000, $90) over ($0, $100), then it might be interpreted as if he is willing to "sacrifice" in order for player 1 to get a larger reward. So, in order for player 2 to figure out for himself whether he is willing at all to sacrifice for player 1, player 2 has to consider the Left sub-game and imagine why player 1 has chosen not to play Left. The Right sub-game yields player 2 more money, but when player 1 chooses Right, it is not clear that he does so out of good will or because of fear and mistrust. If he chooses Left it is not clear whether he does it because he does not believe player 2 would sacrifice for him or that he does not believe player 2 would, in the final account, punish him. Such an analysis could be carried out when the number of types is small as suggested by Harsanyi. But this possibility is due to the very small size of the game. In general, the number of types can be quite large.

It may be argued that the issues of revealed personalities could be addressed by expanding the game to accommodate all possible types of the players. In an extreme situation, each pure strategy may be associated with a different type, so playing a best response relative to a known probability distribution over types seems to solve the problem in principle. Incidentally, in this extreme situation there is no room left for game-theoretic strategic analysis because the prior beliefs about types determine an optimal strategy. It may seem that there is no need to assign more types than strategies, but the problem remains as to how to *interpret* these types. A single strategy can be interpreted in various ways, depending on the types of other players. So, a more appropriate estimate of the number of required types is to assign each player as many types as the number of strategy *tuples*, which

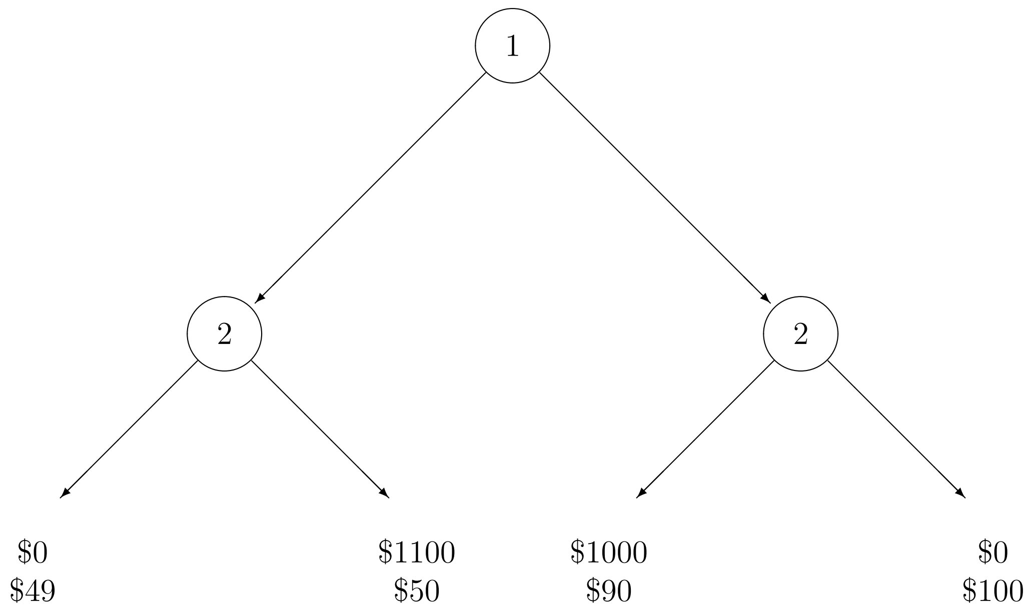

Figure 7: What are the types?

is of course out of the question from a practical point of view. To further explain the difficulties, consider two players where player 1 can be one of n types $T_1^1, \ldots T_n^1$. Now, the type of player 2 has to specify how it relates to each of the possible types of player 1 for example, whether or not he likes that type of 1. This implies that at least 2^n types may be required for player 2. Then, a refinement of the types of player 1 may be needed, and so on. In any case, in order to interpret the different possible kinds of behavior as different personality types, some strategic analysis has to be carried out, and a related question is raised whether or not the determination of the relevant types and the strategic analysis can be de-coupled.

Of course, playing this game repeatedly for a large known finite number of times yields much more problematic examples with many more types.

7 Epilogue: Comparison with psychological games

The subject of this paper seems to be related to the so-called psychological games of Geanakoplos, Pearce and Stacchetti [2]. The issues there, however, are quite different. Psychological games generalize standard games to include beliefs or expectations of players about actions and beliefs or expectations of other players, so utility payoffs depend not only on actions but also on beliefs about actions. Thus, a player my be disappointed if his beliefs about actions of other players do not realize. In this paper, however, we do not interject any additional beliefs into the game. We are interested, from a pragmatical point of view, in the dependence of preferences (over the outcomes) on actions actually taken during the play of the game, as well as ones that could have been taken but were not. Thus, as in the standard theory (and unlike psychological games), in our framework, the utility values of outcomes depend only on the game tree and on the play but not on any external entities. The utility values in our framework, however, are not given as part of the game description but have yet to be determined, in view of some partial preferences and after some analysis of the game. In our opinion, this is a practical challenge. We argue that in some real-life games, where the outcomes are not necessarily given together with utility values, derivation of utility values requires some strategic analysis of the game and cannot be carried out in advance.

The authors of [2] state that "...the traditional theory of games is not well suited to the analysis of such belief-dependent psychological considerations as surprise, confidence..." For our purposes, the framework of utility payoffs depending only on actions (and not on beliefs about actions) is sufficient, but the utility values are not given in advance. Thus, in psychological games beliefs can affect utilities, whereas here, the main cause for revision of utilities is the behavior of players and the effect it has on the attitudes of players toward each other. In our view, if psychological games were to be implemented in practice, similar problems of de-coupling the formation of beliefs and strategic analysis would arise.

References

[1] J.-L. Ferreira, I. Gilboa and M. Maschler, "Credible Equilibria in Games with Utility Changing during the Play," *Games and Economic Behavior* **10** (1995) 284–317.

[2] J. Geanakoplos, D. Pearce and E. Stacchetti, "Psychological games and sequential rationality", *Games and Economic Behavior* **1** (1989) 60–79.

[3] J. Harsanyi, "Games with incomplete information played by Bayesian players," Parts I, II, and III, *Management Science* **14** (1967–8) 159–182, 320–334, 486–502.

[4] G. Owen, *Game Theory*, third edition, Academic Press, New York, 1995.