RJ 10218 (A0109-024) September 24, 2001
Computer Science

IBM Research Report

Formation of Preferences and Strategical Analysis: Can They Be De-Coupled?

Nimrod Megiddo
IBM Research Division
Almaden Research Center
650 Harry Road
San Jose, CA 95120

Research Division
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

This page intentionally left blank.

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

Nimrod Megiddo

September 20, 2001

1 Introduction

In traditional game theory, the analysis of a game relies on individual utility values associated with the various outcomes of the game. In particular, in an extensive form game, each player is assumed to have a utility function defined over the set of leaves of the game tree. In games of incomplete information, such utility values may not be common knowledge. The standard remedy to this situation, following Harsanyi [3], is to expand the game to account for the various possible "types" of players which are jointly drawn from some commonly known probability distribution. The assumption is that such an expansion of the game can be constructed without strategically analyzing the game, since such an analysis would obviously depend on utilities from various outcomes. In this paper, we argue that the formation of preferences over outcomes, and the strategic analysis of the game, may be intertwined to the extent that common resolution procedures, such as backwards induction, cannot not be applied.

The subject of this paper seems to be related to the so-called psychological games of Geankoplus, Pearce and Stacchetti [1]. The issues, however, are quite different. Psychological games include beliefs of players about actions and beliefs of other players, and the utility values depend not only on actions but also on beliefs. Here, we do not introduce any additional beliefs into the game. We are interested in the dependence of preferences (over the outcomes) on actions chosen during the play of the game. Thus, as in the traditional theory (and unlike psychological games), we work with games where the utility values of outcomes depend only on the actions of players. However, the utility values are not given as part of the game but rather have to be determined in view of some given 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 [1] 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 suffices, 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 impact it has on the attitudes of players towards each other.

2 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. This is true, for example, when the game is defined by a succinct set of rules, and the tree encodes all the possible plays. Certainly, the naive approach of determining preferences simply over the set of distinct physical outcomes does not suffice for capturing what happens in real situations. In this section, we present a few examples that can still be handled by the common game theoretic approach. These examples are presented merely for the sake of clarifying the underlying issues. Later, we explain how these examples can be extended into more problematic situations.

Consider, for example, the game depicted in Fig. 1, which is a simplified version of the ultimatum game [2]. 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 (i) zero payoff to either player and (ii) a 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 a second game, where player 1 chooses between (i) zero payoff to either player, and (ii) letting player 2 choose between (\$2,\$2) and (\$3,\$1). Thus, the physical outcomes in these two games are identical. In the first 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 second game. Returning to the first game, the preference of player 2 between (\$0,\$0) and (\$3,\$1) seems to depend on the fact that if indeed he had to choose between the two, then it would be because player 1 had rejected the (\$2,\$2) outcome. That choice does affect how player 2 evaluates the benefit to player 1. He might be so 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) outcome over (\$3,\$1). So, the outcomes that

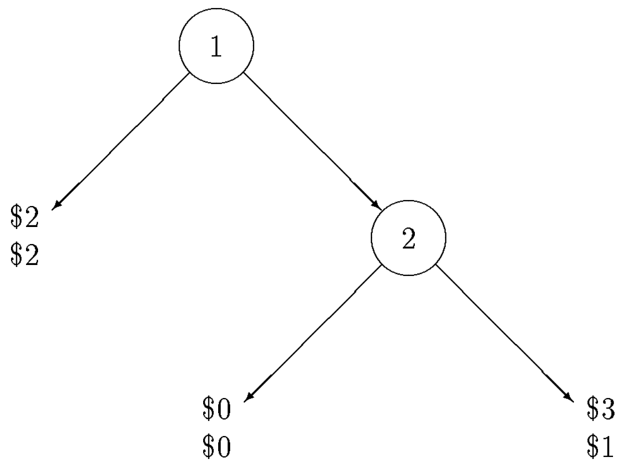

Figure 1: A simplified ultimatum game

could have been reached but have already been rejected, affect the preferences over outcomes that could still be reached.

Within a Bayesian framework, player 1 would have a certain posterior probability p that player 2 would prefer ($0,$0) over ($3,$1), given that player 1 had rejected the ($2,$2) outcome. If $u_1(\cdot,\cdot)$ is player 1's utility function, then he should reject the ($2,$2) outcome if and only if $u_1(2,2) < pu_1(0,0) + (1-p)u_1(3,1)$. The utility of player 1 from each of the outcomes must, however, reflect the context in which the outcome is reached, rather than its mere monetary value. For example, if the outcome ($3,$1) is reached, it must also involve some of anger on the part of player 2, which would impact in some way the utility value for player 1, depending on the latter'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). This preference comes to play when player 1 actually has to choose between these two outcomes. The context in which this choice has to be made can be described as follows. Player 1 has previously chosen not to terminate the game with the outcome of ($10,$0), but rather let player 2 choose. As a consequence, player 2 could have terminated 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 could be seen by player 1 as quite selfish. First, it was player 1 who made it possible for player 2 to receive a nonzero payoff, and player 2 could have secured $100 for each player. But, instead, it seems that player 2 has aimed at gaining an extra single

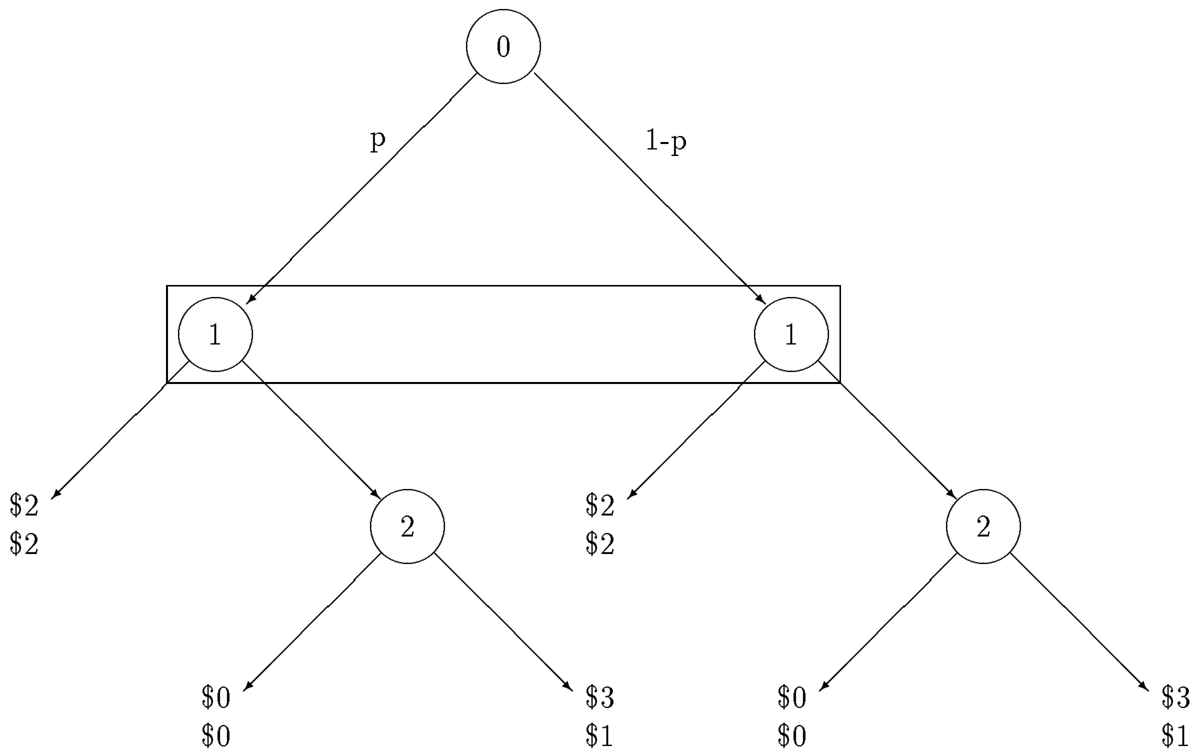

Figure 2: A Bayesian expansion of the simplified ultimatum game

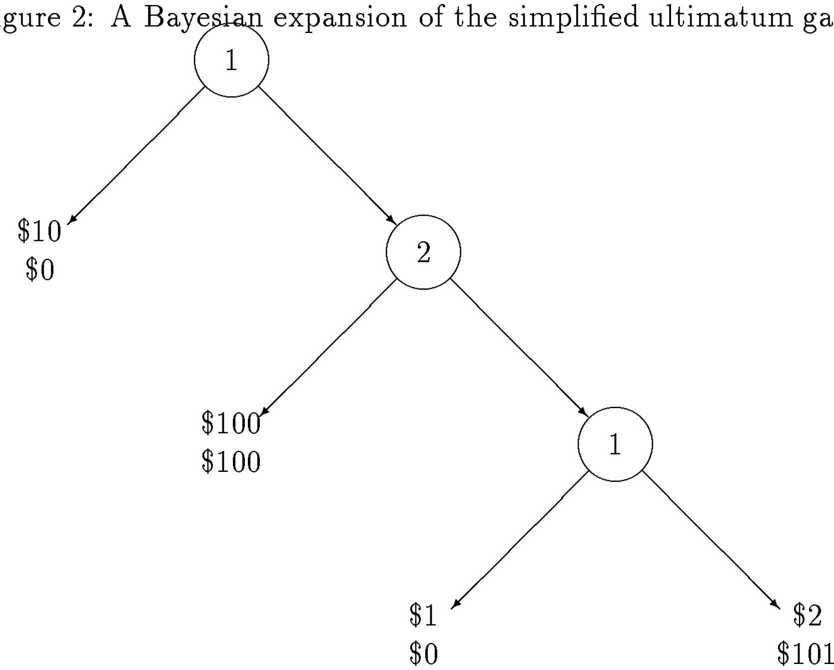

Figure 3: Punishing a greedy player

dollar, at the expense of player 1 receiving only \$2. Thus, player 2 must have assumed that in the final count, player 1 would 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 with that outcome must have been the hope for the \$100 to either player. 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 very specific context of actions taken by both players. It is a consequence of certain interpretations of actions, which cause player 1 to evaluate player 2's personality in a certain way. Here, the argument does *not* rely on beliefs of player 1 about player 2. It is simply the unpleasant personality of player 2, exposed by the latter's actions, that causes player 1 to sacrifice some payoff in order to "teach player 2 a lesson."

Another example of the effect of an "exposed personality" is depicted in Fig. 4. Here, player 2 has to choose between (\$100,\$10,\$0) and (\$200,\$9,\$0), subsequent to player 1 choosing 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 be nice (to player 3), so player 2 might not be nice to him.

3 Simultaneous valuation of disjoint subgames

The main reason, why the examples of the preceding section could still be handled by the traditional game-theoretic approach, is that they do not require simultaneous valuations of disjoint subgames. For example, in the game depicted in Fig. 3, player 2 can terminate the play with payoffs (\$100,\$100) rather than enter a subgame that is anticipated to yield (\$100,\$100). Thus, it is relatively easy to interpret, in that game, a choice of player 2 not to go right, as an expression of greediness and disregard for player 1, whose previous choice to go right benefits 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, etc. 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 decision to go right. On the other hand, if the node (\$100,\$100) were replaced by a certain nontrivial subgame, then in order for player 1 to be able to interpret player 2's decision to avoid such a subgame, he would first have to analyze that nontrivial subgame. Of course, player 2 also would need to carry out such an analysis in order to understand player 1's anticipated reaction. In particular, the subgame could be

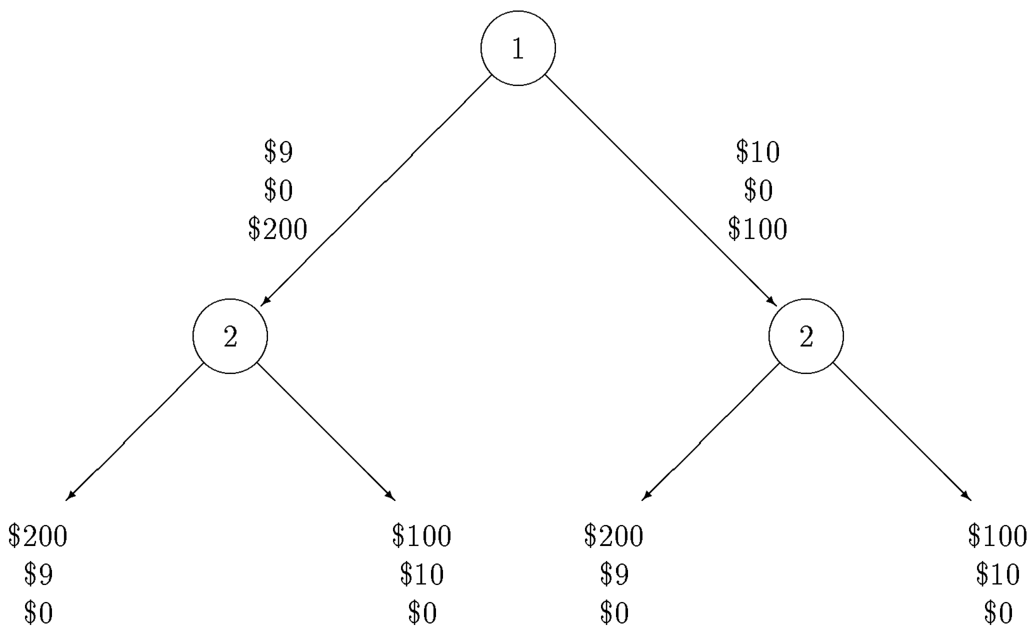

Figure 4: Learning from a player's treatment of a third party

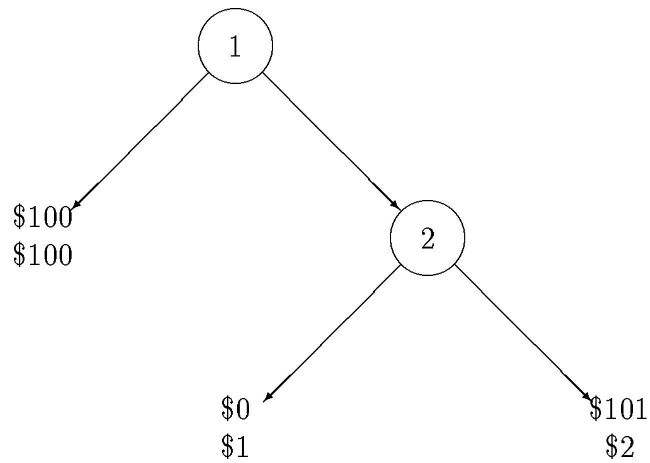

Figure 5: A nontrivial subgame

the one shown in Fig. 5, where the roles of the players are reversed. Thus, in order to valuate this subgame, one would have to appropriately interpret the decision of player 2 to play this subgame rather than let player 1 choose between ($1,$0) and ($2,$101).

At this point, it should be quite clear that, in general, a subgame has to be evaluated in the *context* of how 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 the various outcomes of the subgame. In particular, the valuation of opportunities that are missed by deciding to play a certain subgame affects the valuation of the subgame itself. This fact gives rise to the *circularity* in the analysis. To illustrate this problem, consider a situation, where player 1, in his first move, chooses which subgame, L or R, will be played by himself and player 2 (see Fig. 6), and these subgames are non-trivial. Thus, player 1 has to figure out which subgame he really prefers to play. We argue that player 1 cannot arrive at such a preference by analyzing one of the subgames first and then the other one. Suppose the player 1 attempts to analyze the subgame L first. One important issue that arises in the subgame L is the anticipated reaction of player 2 to the choice of player 1 not to play the subgame R. In order to interpret this choice appropriately, player 2 (and therefore also player 1) would have to analyze the subgame R, in order to see, for example, whether player 1 has acted generously or selfishly. This distinction is important because it affects what player 2 may be willing to do while playing the

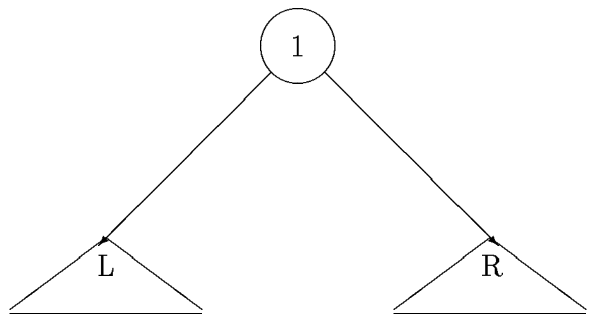

Figure 6: A symmetric circularity

subgame L. Of course, this requirement contradicts the assumption that the subgame L can be analyzed first. It turns out that unless one the subgames is trivial, the valuation of the two subgames has to be carried out *simultaneously*.

The discussion above suggests that, in general, games cannot be analyzed by considering their subgames sequentially. At best, perhaps, one can only achieve consistency of preferences and choices of actions. In the next section, we propose one possible concept of consistency.

4 Consistency of preferences and choice of action

Our point of departure from traditional game theory has to do with *uniqueness* of preferences. We question the assumption that a player must have unique preference order over the outcomes of the game, even before he analyzes the game. In particular, we believe the existence of a unique von Neumann - Morgenstern utility function is not guaranteed. The analysis of the game may give rise to such a preference order but, a priori, it may not be unique. We argue that the strategic choices and the preferences have to be consistent in a sense that we be explain below.

First, the preferences of a player over the set Ou of outcomes are restricted by the player's own fundamental principles. However, the particular utility function, with which consistency of actions is required, is not given in advance. Rather, the utility function $u_i : Ou \to \Re$ of player i must lie in some compact subset U^i of $\Re^{Ou}$, the set of plausible utility functions for player i. If Nash-equilibrium (NE) is adopted as the solution concept, then the solution of a game here would consist of particular utility functions $u^i \in U^i$ and strategies that are in NE with respect to the u^is. But there must be an additional requirement of consistency, namely, the utility functions must be consistent with the actions of the players. In order to satisfy the latter, we propose the concept of a "utility revision map" Ψ as follows. Denote by U and Σ the cartesian products of the plausible utility sets, and the strategy spaces, respectively. The map $\Psi : U \times \Sigma \to U$ models the re-evaluation of outcomes, which players may apply, in view of the choices other players make, and in view of their own previously estimated utilities. Thus, a utility value, which is tentatively assigned to a certain outcome, may be revised after the choices 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.

A utility profile $u \in U$ is said to be a fixed-point relative to a strategy profile $\sigma \in \Sigma$, if $u = \Psi(u, \sigma)$. For a fixed utility profile u, there exists a well-known continuous map $\Phi_u : \Sigma \to \Sigma$, which is used for proving the existence of a Nash-equilibrium (see, for example, []). Thus, for modeling coherence of preferences and actions, we may consider the 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 is obvious. Such a fixed point is a coherent combination of plausible preferences and strategies.

4.1 Combined psychological effects

The difficulties raised above can become even more severe if one is willing to accept the additional complications presented by psychological games. Thus, the interpretation of actions may depend not only strategic analysis but also on beliefs that are added to the game as in [1]. A seemingly simple game can become quite complicated. 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". Suppose that the basic preferences, outside the current game, are that each player would rather have the large object. But within the current game, the other player gets the other object, so the preferences are more subtle. 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 object and who gets the small one, and (ii) let player 2 choose 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 or himself (and leaving the large one for player 1). There could be a number of reasons why player 1 might 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 characteristic 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.

References

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

[2] W. Güth, R. Schmittenberger and B. Schwarze, " An experimental analysis of ulitmatum bargaining," *J. Econ. Behav. Organ.* **3** (1982) 367–388.

[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.