A Computational Characterization of Multiagent Games with Fallacious Rewards

Ariel D. Procaccia   Jeffrey S. Rosenschein
School of Engineering and Computer Science
The Hebrew University of Jerusalem, Israel
{arielpro,jeff}@cs.huji.ac.il

ABSTRACT

Agents engaged in noncooperative interaction may seek to achieve a Nash equilibrium; this requires that agents be aware of others' rewards. Misinformation about rewards leads to a gap between the real interaction model—the *explicit* game—and the game that the agents perceive—the *implicit* game.

If estimation of rewards is based on modeling, agents may err. We define a *robust equilibrium*, which is impervious to slight perturbations, and prove that one can be efficiently pinpointed. We then relax this concept by introducing *persistent equilibrium pairs*—pairs of equilibria of the explicit and implicit games with nearly identical rewards—and resolve associated complexity questions.

Assuming that valuations for different outcomes of the game are *reported* by agents in advance of play, agents may choose to report false rewards in order to improve their eventual payoff. We define the GAME-MANIPULATION (GM) decision problem, and fully characterize the complexity of this problem and some variants.

Categories and Subject Descriptors

F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity;
I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent Systems*

General Terms

Algorithms, Theory

Keywords

Computational complexity, Game theory

1. INTRODUCTION

Game theory has long been a standard tool in the analysis of multiagent interactions. Cooperative game theory, for instance, is widely regarded as the standard model in the context of coalition formation among autonomous agents. Noncooperative game theory is, however, the field in which a large portion of multiagent research is truly grounded, as one often considers settings where agents choose actions in their own self-interest [15].

Nash equilibrium is considered to be the prominent solution concept for noncooperative games; when players' strategies are in Nash equilibrium, no player stands to gain by unilaterally changing its strategy. The applicability of this concept to multiagent systems has motivated researchers to study the complexity of finding a Nash equilibrium, in terms of both the computation required [5, 7] and the communication involved [6].

Unfortunately, it seems impossible to efficiently find a Nash equilibrium, and, to make things worse, agents usually lack complete information about their peers and environment. This has led researchers to devote much effort to learning in games, and in particular learning a Nash equilibrium [9, 10, 1].

Whichever algorithm is used in order to guarantee that agents reach an equilibrium, it is usually assumed that agents can observe others' rewards for given combinations of actions. But what are the possible sources of these observations? We suggest three: *direct observation*—an agent can see another's reward (for instance, in a poker game); *modeling* the other agents and the environment; and *preference revelation*—each agent announces its valuation for different outcomes of the game.

In the first case, it seems reasonable that agents might observe the *correct* rewards for other agents; in the last two cases, this is not necessarily true. Below we use the *implicit games* terminology proposed by Bowling and Veloso [2]: the explicit game is the real game, whereas the implicit game is the same game but with possibly different rewards.[1]

In the overall context of estimating rewards through modeling, we study equilibrium concepts that satisfy some notion of stability in the face of changing rewards. Our idea of an ϵ-robust equilibrium is a strategy profile that remains an equilibrium in any implicit game in which the rewards are perturbed by at most ϵ. We characterize these equilibria and show that they can be found efficiently. Since robust equilibria rarely exist, we produce a relaxation in the form of *persistent equilibrium pairs*, and study the complexity of deciding the existence of such pairs.

The second setting we study is when errors originate in

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.

[1]A somewhat simpler version of Bowling and Veloso's implicit games model is appropriate for our purposes.

strategic preference revelation. The subject of the complexity of manipulation, especially in voting, has been the source of much interest [4, 3, 14]. In our predicament, an agent may improve its payoff by smartly revealing rewards that incite the other agents to achieve an equilibrium of the induced implicit game, but in fact itself plays a strategy that maximizes its own payoff in the explicit game.[2]

In more detail, the other gullible agents, observing the implicit game, believe that the manipulator would play a certain strategy, which is a part of a Nash equilibrium strategy profile in the implicit game; they think this strategy is a best response for the manipulator, and hence would be played. However, the manipulator, aware of its true rewards in the explicit game, plays a different strategy, which improves its outcome. The final outcome the manipulator receives is often better than the one it gets for a strategy profile that is an equilibrium in the explicit game.

We formalize the above discussion in the form of a decision problem—GAME-MANIPULATION (GM)—and prove that the problem is in $\mathcal{P}$ for 2 players, but $\mathcal{NP}$-complete for at least 3 players. We also examine three variants of GM.

2. PRELIMINARIES

We start with a short introduction to game theory, and a formulation of some relevant results.

DEFINITION 1. A *game in strategic form* (hereinafter, simply *game*) G is a tuple $(N, A_1, \ldots, A_n, R_1, \ldots, R_n)$. $N = \{1, \ldots, n\}$ is the set of players, A_i is the finite action set of player i, and $R_i : A \to \mathbb{R}$ is the reward function of player i, where $A = A_1 \times \cdots \times A_n$.

Many of our examples will rely on 2-player games where each player has two actions. We refer to the actions of player 1 as U (up) and D (down), and to the actions of player 2 as L (left) and R (right). If one of the players has 3 actions, we refer to the additional action as M (middle).

DEFINITION 2. A *strategy* for player i is $\pi_i \in \Delta(A_i)$; it specifies the probability of player i playing each action in its action set. A *pure* strategy gives probability 1 to some action. A *mixed* strategy is a strategy that is not pure. A *strategy profile* is a tuple containing one strategy for each player.

For a strategy profile π, we denote by π_{-i} the tuple containing the strategies of all players other than i, and by π_C the tuple of strategies of the players in $C \subseteq N$.

A special case of games is *zero-sum* games. In such games, there are two players with diametrically opposed interests, i.e., $R_1 \equiv -R_2$. When dealing with zero-sum games, we write R instead of R_1; the goal of player 2 is to minimize the reward. It is known that in zero-sum games:

$$\max_{\pi_1 \in \Delta(A_1)} \min_{\pi_2 \in \Delta(A_2)} R(\pi_1, \pi_2) = \min_{\pi_2 \in \Delta(A_2)} \max_{\pi_1 \in \Delta(A_1)} R(\pi_1, \pi_2).$$

This value is denoted by $v(G)$. Strategies that guarantee that the reward be at least $v(G)$ (for player 1) or at most $v(G)$ (for player 2) are called *optimal strategies*.

[2]This is reminiscent of, though not identical to, the strategically lying agents in [15].

LEMMA 1. *[11] In a zero-sum game G, a strategy π_1 for player 1 is optimal iff:*

$$v(G) = \min_{a_2 \in A_2} R(\pi_1, a_2).$$

Similarly, a strategy π_2 for player 2 is optimal iff:

$$v(G) = \max_{a_1 \in A_1} R(a_1, \pi_2).$$

Returning to general-sum games, we discuss the concept of equilibrium. Given a strategy profile $\pi = (\pi_1, \ldots, \pi_n)$, we say that π_i is a *best response* (BR) to π_{-i} iff for all $\pi'_i \in \Delta(A_i)$: $R_i(\pi) \geq R_i(\pi'_i, \pi_{-i})$.

DEFINITION 3. A strategy profile π is a *Nash equilibrium* (NE) iff for all i, π_i is a best response to π_{-i}.

It is a well-known fact that every game has a Nash equilibrium.

LEMMA 2. *[12, Lemma 33.2] π_i is a BR to π_{-i} in G iff every pure strategy in the support of π_i is a BR to π_{-i}.*

REMARK 1. A pair of optimal strategies in a zero-sum game is a Nash equilibrium.

It is still unclear what the complexity of finding a Nash equilibrium (NE) is. Nevertheless, some related questions have been resolved.

THEOREM 1. *[5] Even in symmetric 2-player games, it is $\mathcal{NP}$-hard to decide the following problems:*

1. *Whether there is a NE where player 1 sometimes plays $a_1 \in A_1$.*
2. *Whether there is a NE where all players have expected utility greater than k.*
3. *Whether there is a NE where the social welfare is greater than k.*
4. *Whether there is more than one Nash equilibrium.*

Littman [10] defines a specialized type of equilibrium in which we will be interested.

DEFINITION 4. A Nash equilibrium π is a *coordination* equilibrium iff all players achieve their highest possible value:

$$R_i(\pi) = \max_{a \in A} R_i(a) \tag{1}$$

for all $i \in N$.

3. MODELING INACCURATE REWARDS

In this section, we discuss scenarios where players achieve equilibria based on erroneous models. In such settings, an agent estimates the payoffs other agents receive for action profiles, relying on knowledge of its peers. It is therefore likely that an agent might inaccurately estimate other agents' rewards. A more extraordinary situation is the one where an agent fails to correctly determine its own rewards; this may happen when an agent models the game using an inaccurate model of the environment. The mistakes are often assumed to be bounded by a (relatively small) number ϵ.

Observing the formulation of Bowling and Veloso [2], we dub the "real" game—the game with the actual rewards—the *explicit* game, and the game with the same players and action sets, but perhaps different rewards, the *implicit* game. The explicit game is denoted by

$$G = (n, A_1, \dots, A_n, R_1, \dots, R_n),$$

while the implicit game is represented by

$$\widetilde{G} = (n, A_1, \dots, A_n, \widetilde{R}_1, \dots, \widetilde{R}_n).$$

Definition 5. Let G be an explicit game, and $\widetilde{G}$ be an implicit game. The pair $\langle G, \widetilde{G} \rangle$ is an *ϵ-perturbed system* iff for all $a \in A$ and all $i \in N$: $|R_i(a) - \widetilde{R}_i(a)| \leq \epsilon$.

Example 1. We briefly demonstrate how small changes in the rewards of players in the explicit game significantly impact players' utilities in the Nash equilibrium. In all examples, the explicit game will be specified using parentheses, while the implicit game will be described using square brackets.

$$\begin{pmatrix} (-10,-10) & (9,-9) \\ (-9,9) & (10,10) \end{pmatrix}$$

The game has a single equilibrium at (D, R), with payoffs $(10, 10)$. Several small errors create the implicit game:

$$\begin{bmatrix} (-10,-10) & (9,-11) \\ (-11,9) & (8,8) \end{bmatrix}$$

This game has a single equilibrium at (U, L) with payoffs $(-10, -10)$.

3.1 Robust Equilibria

It is desirable that games be robust to small perturbations of the rewards, in the sense that there is a strategy profile that is an equilibrium in every perturbed implicit game. Should this property hold, mistakes would little affect the outcome of the game.

Definition 6.

1. Let G be a game, and π a strategy profile. We say that π_i is an *ϵ-robust best response* to π_{-i} in G iff for all games $\widetilde{G}$ such that $\langle G, \widetilde{G} \rangle$ is an ϵ-perturbed system, π_i is a best-response to π_{-i} in $\widetilde{G}$.

2. Let G be a game. A strategy profile π is an *ϵ-robust equilibrium* iff for all i, π_i is an ϵ-BR to π_{-i} in G.

Remark 2. Robust equilibrium is clearly a special case of *ex-post* Nash equilibrium in games with incomplete information.

Example 2. In the following explicit game, (U, L) is a 1-robust equilibrium, while (D, R) is an equilibrium that is not ϵ-robust for any $\epsilon > 0$.

$$\begin{pmatrix} (2,2) & (0,0) \\ (0,0) & (0,0) \end{pmatrix}$$

It is often advantageous for agents to jointly achieve a robust equilibrium—if such an equilibrium exists. How hard is it to find one? We first address the existence problem.

Definition 7. In the Robust-NE problem, we are given a game G, a strategy profile π, and $\epsilon > 0$, and are asked whether π is an ϵ-robust NE in G.

Lemma 3. *Let G be a game, and π_i be an ϵ-robust BR to π_{-i} in G. Then π_i is a pure strategy.*

Proof. Assume that the claim is not true. Hence, there exists a game G, and a strategy profile π, s.t. π_i is a BR to π_{-i}, but $\pi_i(\widehat{a}_i) > 0$ and $\pi_i(\widehat{a}'_i) > 0$ for $\widehat{a}_i, \widehat{a}'_i \in A_i$. By Lemma 2, both $\widehat{a}_i$ and $\widehat{a}'_i$ must be BR to π_{-i} in all $\widetilde{G}$ such that $\langle G, \widetilde{G} \rangle$ is an ϵ-perturbed system. Observe the game $\widetilde{G}$, which is identical to G, except that $\widetilde{R}_i(a) = R_i(a) + \epsilon$ for all a such that $a_i = \widehat{a}_i$. As

$$R_i(\widehat{a}_i, \pi_{-i}) = R_i(\widehat{a}'_i, \pi_{-i}),$$

it must hold that

$$\widetilde{R}_i(\widehat{a}_i, \pi_{-i}) = \widetilde{R}_i(\widehat{a}'_i, \pi_{-i}) + \epsilon.$$

This is a contradiction. □

Lemma 4. *$\widehat{a}_i \in A_i$ is an ϵ-robust BR to π_{-i} iff for all $a_i \neq \widehat{a}_i$,*

$$R_i(\widehat{a}_i, \pi_{-i}) \geq R_i(a_i, \pi_{-i}) + 2\epsilon. \quad (2)$$

Proof. Assume first that $\widehat{a}_i$ is an ϵ-robust BR. For some $\widehat{a}'_i \neq \widehat{a}_i$, observe the implicit game in which $\widetilde{R}_i(a) = R_i(a) - \epsilon$ for all $a \in A$ such that $a_i = \widehat{a}_i$, and $\widetilde{R}_i(a) = R_i(a) + \epsilon$ for all $a \in A$ such that $a_i = \widehat{a}'_i$. Clearly it holds that

$$\widetilde{R}_i(\widehat{a}_i, \pi_{-i}) = R_i(\widehat{a}_i, \pi_{-i}) - \epsilon,$$

and

$$\widetilde{R}_i(\widehat{a}'_i, \pi_{-i}) = R_i(\widehat{a}'_i, \pi_{-i}) + \epsilon.$$

As $\widehat{a}_i$ is still a BR in $\widetilde{G}$,

$$\widetilde{R}_i(\widehat{a}_i, \pi_{-i}) \geq \widetilde{R}_i(\widehat{a}'_i, \pi_{-i}).$$

Equation (2) directly follows.

In the other direction, assume Equation (2) holds. We must show that for all $\widehat{a}'_i \in A_i$ and implicit games $\widetilde{G}$ such that $\langle G, \widetilde{G} \rangle$ is an ϵ-perturbed system,

$$\widetilde{R}_i(\widehat{a}_i, \pi_{-i}) \geq \widetilde{R}_i(\widehat{a}'_i, \pi_{-i}). \quad (3)$$

It holds that

$$|\widetilde{R}_i(\widehat{a}_i, \pi_{-i}) - R_i(\widehat{a}_i, \pi_{-i})| \leq \epsilon,$$

and

$$|\widetilde{R}_i(\widehat{a}'_i, \pi_{-i}) - R_i(\widehat{a}'_i, \pi_{-i})| \leq \epsilon.$$

From these equations and our assumption, Equation (3) immediately follows. □

Lemma 5. Robust-NE *is in $\mathcal{P}$.*

Proof. It is sufficient to show that it is possible to determine, in polynomial time, whether a strategy is an ϵ-robust BR. By Lemma 3, we can assume the given strategy π_i is a pure strategy $\widehat{a}_i$. By Lemma 4, we only have to check if Equation (2) holds for all other pure strategies in A_i—and this can be accomplished in polynomial time. □

The theorem follows directly from the above lemmas.

THEOREM 2. *It is possible to find an ϵ-robust equilibrium, or determine that one does not exist, in time that is polynomial in $|A|$.*

REMARK 3. The time required is polynomial in the number of pure action profiles $|A|$, which may be exponential in the number of players n. Nevertheless, notice that our $\mathcal{NP}$-hardness results given below hold even when the size of the input is $|A|$. Moreover, unless a game is special and can be concisely represented, the size of the representation of the game is at least $|A|$.

PROOF OF THEOREM 2. Given a game G, by Lemma 3, in order to determine whether there exists a robust equilibrium it is sufficient to test each pure $a \in A$ for this property. By Lemma 5, the test can be executed in polynomial time. $\square$

REMARK 4. It is possible to define the *robustness level* ϵ^* of a given game as the maximal ϵ such that there exists an ϵ-robust equilibrium in the game. The above results show that the robustness level can also be calculated in polynomial time, using the formula:

$$\epsilon^* = \max_{\widehat{a} \in A} \min_{i \in N} \min_{a_i \neq \widehat{a}_i} \frac{R_i(\widehat{a}) - R_i(a_i, \widehat{a}_{-i})}{2}.$$

Indeed, given $\widehat{a} \in A$, by Lemma 4 the value

$$\min_{i \in N} \min_{a_i \neq \widehat{a}_i} \frac{R_i(\widehat{a}) - R_i(a_i, \widehat{a}_{-i})}{2}$$

is a non-negative number ϵ iff $\widehat{a}$ is an ϵ-robust equilibrium, and is not ϵ'-robust for all $\epsilon' > \epsilon$. By Lemma 3, it is sufficient to maximize over all pure strategy profiles in order to find the robustness level.

3.2 Persistent Equilibrium Pairs

Robust equilibria are characterized by Lemmas 3 and 4; in particular, a robust equilibrium must be a pure strategy profile. It seems that, although efficiently solvable, robust equilibrium is too strong a solution concept. We wish to relax our notion of robustness.

Suppose the players have errors in their game models, and achieve an equilibrium in the implicit game. Even if the new equilibrium is not an equilibrium in the explicit game, we would like to know how "far" (in terms of rewards) this equilibrium is from an equilibrium in the explicit game. If there is a pair of equilibria in the explicit and implicit games that are "close", we think of this pair as satisfying some notion of stability. Of course, the distance between the equilibria depends on the magnitude of the mistakes. Formally:

DEFINITION 8. Let $\langle G, \widetilde{G} \rangle$ be a perturbed system. An ordered pair $\langle \pi, \rho \rangle$ where π is a Nash equilibrium in G and ρ is a Nash equilibrium in $\widetilde{G}$ is an *ϵ-persistent equilibrium pair* of $\langle G, \widetilde{G} \rangle$ iff for all $i \in N$, $|R_i(\pi) - \widetilde{R}_i(\rho)| \leq \epsilon$.

REMARK 5. If $\langle \pi, \rho \rangle$ is an ϵ-persistent equilibrium pair in the ϵ-perturbed system $\langle G, \widetilde{G} \rangle$, then

$$\forall i,\ |R_i(\pi) - R_i(\rho)| \leq 2\epsilon.$$

A persistent equilibrium pair differs from a robust equilibrium in two respects: it is relevant only to a specific perturbed system (instead of every ϵ-perturbed system), and it can hold that $\pi \neq \rho$, as long as the rewards are close.

Before proving some existence theorems for such pairs in special games, we deal with the complexity of determining whether a given perturbed system has a persistent equilibrium pair.

DEFINITION 9. In the PERSISTENT-PAIR problem, we are given an explicit game G, an implicit game $\widetilde{G}$, and a number $\epsilon \geq 0$. We are asked whether $\langle G, \widetilde{G} \rangle$ has an ϵ-persistent equilibrium pair.

THEOREM 3. PERSISTENT-PAIR *is $\mathcal{NP}$-complete, even for two players.*

PROOF. The problem is in $\mathcal{NP}$. Indeed, a specific choice of π and ρ can serve as a witness; it can be determined in polynomial time whether indeed π is a NE in G, ρ a NE in $\widetilde{G}$, and whether their rewards differ by at most ϵ.

We prove that the problem is $\mathcal{NP}$-hard via a reduction from the problem of determining whether a given game has a NE where all players have expected payoff of at least k. Given an instance $\langle G, k \rangle$ of the former problem, the reduction creates an instance of PERSISTENT-PAIR: define $r_{max} = \max_{i \in N, a \in A} R_i(a)$; G is identical to the given game, $\widetilde{G}$ is the game where all rewards are r_{max}, and $\epsilon = r_{max} - k$.[3] Notice that all equilibria in $\widetilde{G}$ have a payoff of r_{max} to all players. Assume that the given instance is a "yes" instance; thus, there is an equilibrium π with payoff of at least k to all players, but clearly with payoff at most r_{max}. Picking some equilibrium ρ in $\widetilde{G}$, we have for all i:

$$|R_i(\pi) - \widetilde{R}_i(\rho)| \leq r_{max} - k = \epsilon.$$

Therefore, $\langle \pi, \rho \rangle$ is an ϵ-persistent equilibrium pair. In the other direction, assume that the given instance is a "no" instance. Any equilibrium π in G has a player i such that $R_i(\pi) < k$. Hence, for any equilibrium ρ in $\widetilde{G}$ it holds that:

$$\widetilde{R}_i(\rho) - R_i(\pi) > r_{max} - k = \epsilon.$$

$\square$

A setting where a persistent equilibrium pair is guaranteed to exist is one where some of the equilibria of the explicit and implicit games satisfy special properties. One such example is requiring that both games have coordination equilibria.[4]

PROPOSITION 4. *Let $\langle G, \widetilde{G} \rangle$ be an ϵ-perturbed system, and let π and ρ be coordination equilibria in G and $\widetilde{G}$, respectively. Then $\langle \pi, \rho \rangle$ is an ϵ-persistent equilibrium pair.*

PROOF. By the definition of an ϵ-perturbed system, it holds that for all $i \in N$ and action profiles $a \in A$:

$$\widetilde{R}_i(a) \geq R_i(a) - \epsilon.$$

As ρ is a coordination equilibrium in $\widetilde{G}$, it follows that for all $i \in N$:

$$\widetilde{R}_i(\rho) = \max_{a \in A} \widetilde{R}_i(a) \geq \widetilde{R}_i(\pi) \geq R_i(\pi) - \epsilon$$

Reversing the roles of R and $\widetilde{R}$, G and $\widetilde{G}$, π and ρ completes the proof. $\square$

[3]It is safe to assume that $r_{max} \geq k$, otherwise the given instance is clearly a "no" instance.

[4]Our interest in coordination equilibria stems from Littman's results: Littman [10] presented a method of achieving Nash equilibrium provided the (stochastic) game has a coordination equilibrium.

Another important case where persistent equilibrium pairs always exist is when G and $\widetilde{G}$ are both zero-sum games.

DEFINITION 10. An ϵ-perturbed system $\langle G, \widetilde{G}\rangle$ is *zero-sum* iff both G and $\widetilde{G}$ are (2 player) zero-sum games.

REMARK 6. In zero-sum perturbed systems, the errors are implicitly assumed to be consistent, in a way: an error e in the reward of one of the players entails an error $-e$ in the reward of the other.

PROPOSITION 5. *Let $\langle G, \widetilde{G}\rangle$ be a zero-sum ϵ-perturbed system, and let π and ρ be optimal strategy profiles in G and $\widetilde{G}$, respectively. Then $\langle \pi, \rho\rangle$ is an ϵ-persistent equilibrium pair.*

PROOF. We have by Lemma 1:

$$v(G) = \min_{a_2\in A_2} R(\pi_1, a_2),$$

and

$$v(\widetilde{G}) = \min_{a_2\in A_2} \widetilde{R}(\rho_1, a_2).$$

Consequently,

$$\begin{aligned}
R(\pi_1,\pi_2) &= v(G)\\
&= \max_{x\in\Delta(A_1)} \min_{a_2\in A_2} R(x, a_2)\\
&\geq \min_{a_2\in A_2} R(\rho_1, a_2)\\
&= R(\rho_1, \widehat{a_2})\\
&= \sum_{a_1\in A_1} \rho_1(a_1) R(a_1, \widehat{a_2})\\
&= \sum_{a_1\in A_1} \rho_1(a_1) \widetilde{R}(a_1, \widehat{a_2})\\
&\quad + \sum_{a_1\in A_1} \rho_1(a_1)\left[R(a_1, \widehat{a_2}) - \widetilde{R}(a_1, \widehat{a_2})\right]\\
&\geq \widetilde{R}(\rho_1, \widehat{a_2}) - \epsilon \sum_{a_1\in A_1} \rho_1(a_1)\\
&= \widetilde{R}(\rho_1, \widehat{a_2}) - \epsilon\\
&\geq \min_{a_2\in A_2} \widetilde{R}(\rho_1, a_2) - \epsilon\\
&= v(\widetilde{G}) - \epsilon\\
&= \widetilde{R}(\rho_1, \rho_2) - \epsilon.
\end{aligned}$$

Reversing the roles of G and $\widetilde{G}$, R and $\widetilde{R}$, π and ρ completes the proof. $\square$

As interesting as persistent equilibria are, their weakness is in the implicit assumption that all players make the same mistakes in the estimation of the rewards—all players observe the same implicit game. What happens if only one of the players errs? In this case, this player (say player 2) converges to an equilibrium strategy ρ in $\widetilde{G}$, while the other player converges to an equilibrium strategy π in G. We would like to know what happens when π_1 is played against ρ_2 in the explicit zero-sum game.

PROPOSITION 6. *Let $\langle G, \widetilde{G}\rangle$ be a 2-player zero-sum ϵ perturbed system, π an optimal strategy profile in G, and ρ an optimal strategy profile in $\widetilde{G}$. Then:*

1. $R(\rho_1, \pi_2) \geq v(G) - 2\epsilon$.
2. $R(\pi_1, \rho_2) \leq v(G) + 2\epsilon$.

PROOF. For part 1, assume the contrary, namely that $R(\rho_1, \pi_2) < v(G) - 2\epsilon$. Then:

$$\begin{aligned}
R(\rho_1, \rho_2) &= v(\widetilde{G})\\
&\leq \widetilde{R}(\rho_1, \pi_2)\\
&\leq R(\rho_1, \pi_2) + \epsilon\\
&< v(G) - \epsilon\\
&= R(\pi_1, \pi_2) - \epsilon.
\end{aligned}$$

This is a contradiction to Proposition 5. The proof of the second part is symmetrical. $\square$

4. REVEALING FALSE REWARDS

In the previous section we examined settings where agents achieve a false equilibrium, as a result of erroneous estimation of rewards. In this section, we look at settings where agents reveal their preferences to other agents, by reporting their valuations for different action profiles. In such cases, an agent may improve its payoff by reporting false valuations. When this happens, we say the lying agent reveals its preferences *strategically*, and refer to it as a *manipulator*.

EXAMPLE 3. The following example presents a successful manipulation by player 1. Observe the explicit game:

$$\begin{pmatrix} (0,0) & (1,1)\\ (1,1) & (2,0)\end{pmatrix}$$

The unique NE is (D, L), with payoff 1 to both players. We would like to know if player 1 can do better. By changing player 1's rewards, we obtain the implicit game:

$$\begin{bmatrix} (0,0) & (1,1)\\ (-1,1) & (-1,0)\end{bmatrix}$$

Observing the implicit game, player 2 achieves the only NE: (U, R), i.e., plays strategy R. Player 2 believes (U, R) is a Nash equilibrium, and thus player 1 can do no better than play U. However, being aware of the real rewards and knowing that player 2 would play R, player 1 counters with D, receiving a payoff of 2 (and leaving player 2 with a sucker's payoff of 0).

The next definition formalizes the above discussion.

DEFINITION 11. In the GAME-MANIPULATION (GM) problem, we are given an explicit n-player game G, a player $\underline{i} \in N$, and an integer $k \in \mathbb{Z}$. We are asked whether there is an implicit game $\widetilde{G}$ where only the rewards of $\underline{i}$ are changed, a (possibly mixed) strategy profile ρ and $\pi_{\underline{i}} \in \Delta(A_{\underline{i}})$, such that ρ is a Nash equilibrium in $\widetilde{G}$, and $R_{\underline{i}}(\pi_{\underline{i}}, \rho_{-\underline{i}}) > k$.

EXAMPLE 4. We construct a "no" instance of the GM problem. The explicit game is:

$$\begin{pmatrix} (0,0) & (1,1) & (0,10)\\ (0,10) & (1,1) & (0,0)\end{pmatrix}$$

Player 1's payoff in any Nash equilibrium is 0; set $\underline{i} = 1$ and $k = 0$. Player 1 cannot improve its payoff by manipulation, since in every Nash equilibrium in the implicit game it holds that player 2 never plays M. Indeed, for any strategy profile, suppose w.l.o.g. that player 1 plays U with probability at least $\frac{1}{2}$. Then player 2's reward for playing action M with probability p is p, but shifting to a strategy that increases the probability of playing R by p increases the reward for player 2 by $5p$.

4.1 Complexity of GM

Our goal in this subsection is to prove:

THEOREM 7 (DICHOTOMY OF GM). *GM with at least 3 players is $\mathcal{NP}$-complete, while GM with 2 players is in $\mathcal{P}$.*

The proof is naturally decomposed into two parts.

PROPOSITION 8. *GM with at least 3 players is $\mathcal{NP}$ complete.*

PROOF. To show that GM is in $\mathcal{NP}$, observe that for a given instance of the problem, a specific example of $\widetilde{G}$, ρ and $\pi_{\underline{i}}$ is a witness; it can be ascertained in polynomial time whether indeed ρ is a Nash equilibrium in $\widetilde{G}$ and $R_{\underline{i}}(\pi_{\underline{i}}, \rho_{-\underline{i}}) > k$.

For the $\mathcal{NP}$-hardness, we prove that the problem of determining whether there exists a NE in a 2-person game where player 1 sometimes plays $\widehat{a}_1 \in A_1$ reduces to GM with 3 players. Given an instance of the former problem, construct as an instance of the latter a three-player game G, where player 3 has only one action. The rewards for players 1 and 2 for any action profile are the same as in the given game (when the action of player 3 is disregarded), and the payoff for player 3 is 1 for any action profile where player 1 plays $\widehat{a}_1$, and 0 otherwise. We also set $k = 0$, and the manipulator to be player 3 ($\underline{i}$=3).

Assume the given instance has a NE π where player 1 plays $\widehat{a}_1$ with probability $r > 0$. Construct the instance of GM as above, and choose $\widetilde{G} \equiv G$. Clearly, when players 1 and 2 use π_1 and π_2 in G (and player 3 plays his single action), this is a Nash equilibrium with the property that player 3 has payoff r.

On the other hand, it is clear that if (π_1, π_2) is not a Nash equilibrium in the given instance, then (π_1, π_2, a_3) (where a_3 is the single action in A_3) is not a Nash equilibrium in any $\widetilde{G}$ where only the rewards of player 3 have changed. Therefore, if the given instance has no NE where player 1 sometimes plays $\widehat{a}_1$, every NE in every admissible $\widetilde{G}$ has the property that player 1 plays $\widehat{a}_1$ with probability 0, and thus the utility of player 3 is 0. $\square$

In the second part of the theorem's proof we require the following lemma. For ease of exposition, in the lemma and subsequent proposition we consider player 1 to be the manipulator, but this is of course an arbitrary choice.

LEMMA 6. *A given instance of GM with two players where player 1 is the manipulator is a "yes" instance iff there exists a pure strategy profile $(\widehat{a}_1, \widehat{a}_2)$ such that $R_1(\widehat{a}_1, \widehat{a}_2) > k$, and a strategy $\rho_1 \in \Delta(A_1)$ such that $\widehat{a}_2$ is a best-response to ρ_1 (in the explicit game).*

PROOF. Assume first that there exists a pure strategy profile $(\widehat{a}_1, \widehat{a}_2)$ such that $R_1(\widehat{a}_1, \widehat{a}_2) > k$, and a strategy $\rho_1 \in \Delta(A_1)$ such that $\widehat{a}_2$ is a best-response to ρ_1. Define $\widetilde{G}$ to be the implicit game where the rewards of player 1 are all identical. It follows from the assumption that $\widehat{a}_2$ is a best-response to ρ_1 and from the construction of $\widetilde{G}$ that $(\rho_1, \widehat{a}_2)$ is a Nash equilibrium in $\widetilde{G}$. Setting ρ_2 from the definition of GM to be $\widehat{a}_2$ and π_1 to be $\widehat{a}_1$ proves that the given instance is a "yes" instance of the GM problem.

In the other direction, assume we are given a "yes" instance of GM (with two players, where player 1 is the manipulator). Let $\widetilde{G}$ be an implicit game where only the rewards of player i are changed, and let $\rho = (\rho_1, \rho_2)$ and $\pi_{\underline{i}} \in \Delta(A_{\underline{i}})$ be strategies such that ρ is a Nash equilibrium in $\widetilde{G}$, and $R_1(\pi_1, \rho_2) > k$. There must be $\widehat{a}_1 \in A_1$ such that $\pi_1(\widehat{a}_1) > 0$ and $R_1(\widehat{a}_1, \rho_2) > k$. Hence, there must be $\widehat{a}_2 \in A_2$ such that $\rho_2(\widehat{a}_2) > 0$ and $R_1(\widehat{a}_1, \widehat{a}_2) > k$. By the assumption, (ρ_1, ρ_2) is a Nash equilibrium in $\widetilde{G}$; from Lemma 2 it follows that any action which is sometimes played in ρ_2 is a best response to ρ_1, and in particular $\widehat{a}_2$ is a best-response to ρ_1. $\square$

PROPOSITION 9. *GM with 2 players is in $\mathcal{P}$.*

PROOF. We present an algorithm for GM, and prove that it always terminates in polynomial time.

```
1: for each (â1, â2) ∈ A s.t. R1(â1, â2) > k do
2:     if ∃ρ1 ∈ Δ(A1) s.t. â2 is a BR to ρ1 then
3:         return true
4:     end if
5: end for
6: return false
```

It follows straightforwardly from Lemma 6 that a given instance of GM is a "yes" instance iff the above algorithm accepts. Therefore, we only need to show that the algorithm can be implemented efficiently.

The algorithm performs the test in line 2 a polynomial number of repetitions; this test can also be performed in polynomial time. Indeed, there is ρ_1 such that $\widehat{a}_2$ is a best response to ρ_1 iff there is a feasible solution to the linear program with the following constraints:

$$\textbf{find } \rho_1(a_1) \textbf{ such that}$$

$$\forall a_2 \in A_2, \sum_{a_1 \in A_1} \rho_1(a_1)[R_2(a_1, \widehat{a}_2) - R_2(a_1, a_2)] \geq 0$$

$$\sum_{a_1 \in A_1} \rho_1(a_1) = 1$$

$$\forall a_1 \in A_1,\ 0 \leq \rho_1(a_1) \leq 1$$

It is well-known that it is possible to determine whether a linear program has a feasible solution in polynomial time. This completes the proof of the proposition, as well as the proof of Theorem 7. $\square$

4.2 Variations on GM

In this subsection, we are concerned with variations on the GM problem.

GM is naturally generalized to a setting where there are several manipulators. The manipulators are all interested in securing a payoff greater than k. However, each manipulator is still selfish, and does not hesitate to improve its own payoff at the expense of other manipulators. Hence, we require that the manipulators' strategies in the explicit game also be in equilibrium.

DEFINITION 12. In the COALITIONAL-GAME-MANIPULATION (CGM) problem, we are given an explicit n-player game G, a subset of players $C \subseteq N$, and an integer $k \in \mathbb{Z}$. We are asked whether there is an implicit game $\widetilde{G}$ where only the rewards of the players in C are changed, a (possibly mixed) strategy profile ρ and $\pi_C \in \prod_{i\in C} A_i$, such that ρ is a Nash equilibrium in $\widetilde{G}$, π_i is a best response to (π_{C-i}, ρ_{-C}) in G for all $i \in C$, and $\forall i \in C, R_i(\pi_C, \rho_{-C}) > k$.

EXAMPLE 5. We clarify the definition by showing a coalitional manipulation in a three-player game, in which players 1 and 2 are manipulators. The game is described by two matrices: the left is associated with an action of player 3 that we dub B (backward), while the right matrix is associated with F (forward).

$$\begin{pmatrix} (2,2,0) & (1,1,0) \\ (1,1,1) & (0,0,1) \end{pmatrix} \qquad \begin{pmatrix} (1,1,2) & (0,0,0) \\ (0,0,0) & (0,0,0) \end{pmatrix}$$

In this game there is a Nash equilibrium at (U, L, F), with payoff $(1, 1, 2)$. Now observe the implicit game:

$$\begin{bmatrix} (2,2,0) & (1,1,0) \\ (1,1,1) & (2,2,1) \end{bmatrix} \qquad \begin{bmatrix} (1,1,2) & (0,0,0) \\ (2,2,0) & (0,0,0) \end{bmatrix}$$

(U, L, F) is no longer an equilibrium, but (D, R, B) is. Player 3 plays B, and the manipulators counter with (U, L); the utility of both manipulators is 2. The manipulators are not motivated to deviate, since U is a best response for player 1 against (L, B) in the explicit game, and L is a best response against (U, B).

PROPOSITION 10. *CGM with at least 3 players is $\mathcal{NP}$-complete, while CGM with 2 players is in $\mathcal{P}$.*

PROOF SKETCH. With 2 players, the CGM problem is identical to the GM problem, so the result for this case follows from Proposition 9. For at least 3 players, we examine two cases.[5] If there are at least two non-manipulators, the proof is similar to the proof of Proposition 8; again, we use a reduction from the problem of determining whether a given game has a Nash equilibrium where player 1 sometimes uses a given strategy. In this case, a given instance is reduced to a game where all the manipulators are players with only one action, and their rewards are greater than 0 only for action profiles that include the given action.

If there are at least two manipulators, the proof is just as straightforward, using a reduction from the problem of determining whether a given game has a Nash equilibrium where all players have expected utility at least k. A given two-player game is reduced to an instance where the two manipulators have exactly the same rewards as in the given game, and a single non-manipulator has a single action. $\square$

REMARK 7. It is also possible to formulate the CGM problem for a cooperative setting, where the manipulators act as a coalition. In this case, we would only require that the total reward of the coalition be greater than k, and would drop the requirement that the manipulators' strategies be in equilibrium in the explicit game—but this is outside the scope of this paper.

In the GM problem, the manipulator was only concerned about its own reward. Nevertheless, some "Robin Hood" type manipulators might wish to engage in deceit in order to improve the welfare of all players.

EXAMPLE 6. Consider the explicit game:

$$\begin{pmatrix} (1,1) & (0,0) \\ (0,0) & (-1,10) \\ (0,10) & (9,9) \end{pmatrix}$$

[5]We assume there is at least one non-manipulator.

The unique NE is (U, L), with payoff 1 to both players. In this case, player 1 can significantly increase the payoff to both players by reporting false rewards.

$$\begin{bmatrix} (0,1) & (0,0) \\ (1,0) & (1,10) \\ (0,10) & (0,9) \end{bmatrix}$$

In the implicit game, there is a single equilibrium in (M, R). Player 2 plays R, but player 1 counters with D. Consequently, both players get a payoff of 9. The social welfare increased from 2 to 18.

DEFINITION 13. In the BENEVOLENT-GAME-MANIPULATION (BGM) problem, we are given an explicit n-player game G, a player $\underline{i} \in N$, and an integer $k \in \mathbb{Z}$. We are asked whether there is an implicit game $\widetilde{G}$ where only the rewards of $\underline{i}$ are changed, a (possibly mixed) strategy profile ρ and $\pi_{\underline{i}} \in \Delta(A_{\underline{i}})$, such that ρ is a Nash equilibrium in $\widetilde{G}$, and $\sum_{i\in N} R_{\underline{i}}(\pi_{\underline{i}}, \rho_{-\underline{i}}) > k$.

PROPOSITION 11. *BGM with at least 3 players is $\mathcal{NP}$-complete, while BGM with 2 players is in $\mathcal{P}$.*

PROOF SKETCH. BGM with at least 3 players is clearly in $\mathcal{NP}$. For the $\mathcal{NP}$-hardness, we prove that the problem of determining whether there exists a NE in a 2-person game where the expected social welfare is at least k reduces to BGM with 3 players. Given an instance of the former problem, construct as an instance of the latter a three-player game G, where player 3 is the manipulator and has only one action; the rewards for players 1 and 2 for any action profile are the same as before (when the action of player 3 is disregarded), and the payoff for player 3 is 0 for all action profiles. The parameter k of the BGM instance is the same as in the given instance. It is easily verified that this is a polynomial-time reduction.

For the two player setting, we notice that Lemma 6 can be reformulated so as to be useful here, as well; therefore, similarly to the proof of Proposition 9, we need only test each $(\widehat{a}_1, \widehat{a}_2) \in A$ such that $R_1(\widehat{a}_1, \widehat{a}_2) + R_2(\widehat{a}_1, \widehat{a}_2) > k$ for the property: there exists $\rho_1 \in \Delta(A_1)$ such that $\widehat{a}_2$ is a BR to ρ_1. This, as before, can be done in polynomial time. $\square$

Returning to the selfish manipulator setting (the GM formulation), another issue requires attention. Even if the manipulator succeeds in finding $\widetilde{G}$, ρ, and $\pi_{\underline{i}}$ as before, it is not at all certain that the other players, observing the implicit game $\widetilde{G}$, would achieve ρ. This is only guaranteed if ρ is a unique NE in $\widetilde{G}$. We define a strong manipulation to be exactly as before, except that we now require that ρ be unique.

DEFINITION 14. In the STRONG-GAME-MANIPULATION (SGM) problem, we are given an explicit n-player game G, a player i, and an integer k. We are asked whether there is an implicit game $\widetilde{G}$ where only the rewards of player i are changed, a (possibly mixed) strategy profile ρ and $\pi_{\underline{i}} \in \Delta(A_{\underline{i}})$, such that ρ is a *unique* Nash equilibrium in $\widetilde{G}$, and $R_{\underline{i}}(\pi_{\underline{i}}, \rho_{-\underline{i}}) > k$.

REMARK 8. Example 3 is in fact an example of a strong manipulation.

PROPOSITION 12. *SGM with at least 3 players is co$\mathcal{NP}$-hard.*

PROOF SKETCH. The complement of the problem of determining whether there exists more than one Nash equilibrium is the problem of determining whether a given game has a unique Nash equilibrium; this problem (with 2 players) is co$\mathcal{NP}$-hard, and it reduces to SGM with 3 players. Indeed, almost the same reduction as in the proof of Proposition 11 also works in this case: the rewards for players 1 and 2 for any action profile are the same as in the given instance (when the action of player 3 is disregarded), and the payoff for player 3 is 0 for all action profiles. k is set to be $-\infty$. Again, it can be verified that this is a reduction. □

5. CONCLUSIONS

We have demonstrated that when agents achieve a Nash equilibrium on the basis of fallacious rewards, their utilities may change substantially. For the setting where agents estimate rewards by relying on a model of other agents and the environment, we have defined the concept of ϵ-robust equilibrium, and have shown that if one exists, it can be found efficiently. However, our characterization of these equilibria implies that they rarely exist. Accordingly, we have relaxed the definition to obtain persistent equilibrium pairs. Although such pairs always exist when the explicit and implicit games are both zero-sum or both have coordination equilibria, in general deciding the existence of such a pair is $\mathcal{NP}$-hard. We wish to comment that the problem of determining whether some multiagent setting has a robust equilibrium or persistent pair is not associated with the agents (which are not aware of the explicit game), but rather with the system designer, who strives to create a stable system.

We have also considered mistakes that are grounded in false reports by manipulative agents. We have shown that the GAME-MANIPULATION problem is in $\mathcal{P}$ for 2 players, and is $\mathcal{NP}$-complete for at least 3 players. In addition, we have demonstrated that similar results hold for the "coalitional", "benevolent", and "strong" variations of GM. These results suggest that manipulation may not be a major concern when there are at least three agents, but it remains to determine how hard it is to decide these problems in the average-case. The "benevolent" setting seems especially interesting: although, in general, one wishes to avoid manipulations, a well-intentioned lie can help agents avoid paying the price of anarchy [13].

It is important to note that the results in this paper are general, in the sense that they are independent of the specific algorithms the agents use to reach an equilibrium. It suffices to assume that they all use such an algorithm (not necessarily the same one). Moreover, the results are stated with respect to games in normal form; this makes them applicable in repeated games, as a Nash equilibrium of the stage game is also an equilibrium of the repeated game when played repeatedly. Nevertheless, multiagent interactions are often modeled in the wider framework of stochastic (Markov) games. A very interesting direction for future research would be to generalize our results to such games.

Another way to extend our research is to derive results concerning other equilibrium concepts, such as correlated equilibria [12]. We are motivated to explore this solution concept specifically, since there has been work on learning correlated equilibria [8]. In the context of manipulation, the problems GM, BGM, and SGM can be defined exactly as before, with the exception that the equilibria are correlated instead of Nash. Problems that are reformulated for correlated equilibria are expected to be easier, as correlated equilibria can be found in polynomial time.

6. ACKNOWLEDGMENT

This work was partially supported by Israel Science Foundation grant #898/05.

7. REFERENCES

[1] M. Bowling and M. Veloso. Multiagent learning using a variable learning rate. *Artificial Intelligence*, 136:215–250, 2002.

[2] M. Bowling and M. Veloso. Existence of multiagent equilibria with limited agents. *Journal of Artificial Intelligence Research*, 22:353–384, 2004.

[3] V. Conitzer, J. Lang, and T. Sandholm. How many candidates are needed to make elections hard to manipulate? In *Proceedings of the International Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 201–214, 2003.

[4] V. Conitzer and T. Sandholm. Complexity of manipulating elections with few candidates. In *Proceedings of the National Conference on Artificial Intelligence*, pages 314–319, 2002.

[5] V. Conitzer and T. Sandholm. Complexity results about Nash equilibria. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 765–771, 2003.

[6] V. Conitzer and T. Sandholm. Communication complexity as a lower bound for learning in games. In *Proceedings of the 21st International Conference on Machine Learning*, pages 185–192, 2004.

[7] I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, pages 80–93, 1989.

[8] A. Greenwald and K. Hall. Correlated Q-learning. In *Proceedings of the 20th International Conference on Machine Learning*, pages 242–249, 2003.

[9] J. Hu and M. Wellman. Nash Q-learning for general-sum stochastic games. *Journal of Machine Learning Research*, 4:1039–1069, 2004.

[10] M. L. Littman. Friend-or-foe Q-learning in general-sum games. In *Proceedings of the 18th International Conference on Machine Learning*, pages 322–328, 2001.

[11] P. Morris. *Introduction to Game Theory*. Springer, 1994.

[12] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[13] C. Papadimitriou. Algorithms, games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing*, pages 749–753, 2001.

[14] A. D. Procaccia and J. S. Rosenschein. Junta distributions and the average-case complexity of manipulating elections. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 497–504, 2006.

[15] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers*. MIT Press, Cambridge, Massachusetts, 1994.