

Computing the Optimal Strategy to Commit to*

Vincent Conitzer
Carnegie Mellon University
Computer Science Department
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
conitzer@cs.cmu.edu

Tuomas Sandholm
Carnegie Mellon University
Computer Science Department
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
sandholm@cs.cmu.edu

ABSTRACT

In multiagent systems, strategic settings are often analyzed under the assumption that the players choose their strategies simultaneously. However, this model is not always realistic. In many settings, one player is able to *commit* to a strategy before the other player makes a decision. Such models are synonymously referred to as *leadership*, *commitment*, or *Stackelberg* models, and optimal play in such models is often significantly different from optimal play in the model where strategies are selected simultaneously.

The recent surge in interest in *computing* game-theoretic solutions has so far ignored leadership models (with the exception of the interest in *mechanism design*, where the designer is implicitly in a leadership position). In this paper, we study how to compute optimal strategies to commit to under both commitment to pure strategies and commitment to mixed strategies, in both normal-form and Bayesian games. We give both positive results (efficient algorithms) and negative results (NP-hardness results).

Categories and Subject Descriptors

J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*; I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems; F.2 [Theory of Computation]: Analysis of Algorithms and Problem Complexity

General Terms

Algorithms, Economics, Theory

Keywords

Game theory, commitment, leadership, Stackelberg, normal-form games, Bayesian games, Nash equilibrium

*This material is based upon work supported by the National Science Foundation under ITR grants IIS-0121678 and IIS-0427858, a Sloan Fellowship, and an IBM Ph.D. Fellowship.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'06, June 11–15, 2006, Ann Arbor, Michigan, USA.
Copyright 2006 ACM 1-59593-236-4/06/0006 ...\$5.00.

1. INTRODUCTION

In multiagent systems with self-interested agents (including most economic settings), the optimal action for one agent to take depends on the actions that the other agents take. To analyze how an agent should behave in such settings, the tools of *game theory* need to be applied. Typically, when a strategic setting is modeled in the framework of game theory, it is assumed that players choose their strategies simultaneously. This is especially true when the setting is modeled as a *normal-form* game, which only specifies each agent's utility as a function of the vector of strategies that the agents choose, and does not provide any information on the order in which agents make their decisions and what the agents observe about earlier decisions by other agents. Given that the game is modeled in normal form, it is typically analyzed using the concept of *Nash equilibrium*. A Nash equilibrium specifies a strategy for each player, such that no player has an incentive to individually deviate from this profile of strategies. (Typically, the strategies are allowed to be *mixed*, that is, probability distributions over the original (*pure*) strategies.) A (mixed-strategy) Nash equilibrium is guaranteed to exist in finite games [18], but one problem is that there may be multiple Nash equilibria. This leads to the *equilibrium selection problem* of how an agent can know which strategy to play if it does not know which equilibrium is to be played.

When the setting is modeled as an *extensive-form game*, it is possible to specify that some players receive some information about actions taken by others earlier in the game before deciding on their action. Nevertheless, in general, the players do not know everything that happened earlier in the game. Because of this, these games are typically still analyzed using an equilibrium concept, where one specifies a mixed strategy for each player, and requires that each player's strategy is a best response to the others' strategies. (Typically an additional constraint on the strategies is now imposed to ensure that players do not play in a way that is irrational with respect to the information that they have received so far. This leads to refinements of Nash equilibrium such as *subgame perfect* and *sequential* equilibrium.)

However, in many real-world settings, strategies are not selected in such a simultaneous manner. Oftentimes, one player (the *leader*) is able to *commit* to a strategy before another player (the *follower*). This can be due to a variety of reasons. For example, one of the players may arrive at the site at which the game is to be played before another agent (*e.g.*, in economic settings, one player may enter a market earlier and commit to a way of doing busi-

ness). Such commitment power has a profound impact on how the game should be played. For example, the leader may be best off playing a strategy that is dominated in the normal-form representation of the game. Perhaps the earliest and best-known example of the effect of commitment is that by von Stackelberg [25], who showed that, in Cournot’s duopoly model [5], if one firm is able to commit to a production quantity first, that firm will do much better than in the simultaneous-move (Nash) solution. In general, if commitment to *mixed strategies* is possible, then (under minor assumptions) it never hurts, and often helps, to commit to a strategy [26]. Being forced to commit to a pure strategy sometimes helps, and sometimes hurts (for example, committing to a pure strategy in rock-paper-scissors before the other player’s decision will naturally result in a loss). In this paper, we will assume commitment is always forced; if it is not, the player who has the choice of whether to commit can simply compare the commitment outcome to the non-commitment (simultaneous-move) outcome.

Models of leadership are especially important in settings with multiple self-interested *software* agents. Once the code for an agent (or for a team of agents) is finalized and the agent is deployed, the agent is committed to playing the (possibly randomized) strategy that the code prescribes. Thus, as long as one can credibly show that one cannot change the code later, the code serves as a commitment device. This holds true for recreational tournaments among agents (*e.g.*, poker tournaments, RoboSoccer), and for industrial applications such as sensor webs.

Finally, there is also an implicit leadership situation in the field of *mechanism design*, in which one player (the designer) gets to choose the rules of the game that the remaining players then play. Mechanism design is an extremely important topic to the EC community: the papers published on mechanism design in recent EC conferences are too numerous to cite. Indeed, the mechanism designer may benefit from committing to a choice that, if the (remaining) agents’ actions were fixed, would be suboptimal. For example, in a (first-price) auction, the seller may wish to set a positive (*artificial*) *reserve price* for the item, below which the item will not be sold—even if the seller values the item at 0. In hindsight (after the bids have come in), this (naïvely) appears suboptimal: if a bid exceeding the reserve price came in, the reserve price had no effect, and if no such bid came in, the seller would have been better off accepting a lower bid. Of course, the reason for setting the reserve price is that it incentivizes the bidders to bid higher, and because of this, setting artificial reserve prices can actually increase expected revenue to the seller.

A significant amount of research has recently been devoted to the computation of solutions according to various solution concepts for settings in which the agents choose their strategies simultaneously, such as dominance [7, 11, 3] and (especially) Nash equilibrium [8, 21, 16, 15, 2, 22, 23, 4]. However, the computation of the optimal strategy to *commit* to in a leadership situation has gone ignored. Theoretically, leadership situations can simply be thought of as an extensive-form game in which one player chooses a strategy (for the original game) first. The number of strategies in this extensive-form game, however, can be exceedingly large. For example, if the leader is able to commit to a mixed strategy in the original game, then every one of the (continuum of) mixed strategies constitutes a pure strategy

in the extensive-form representation of the leadership situation. (We note that a commitment to a distribution is not the same as a distribution over commitments.) Moreover, if the original game is itself an extensive-form game, the number of strategies in the extensive-form representation of the leadership situation (which is a different extensive-form game) becomes even larger. Because of this, it is usually not computationally feasible to simply transform the original game into the extensive-form representation of the leadership situation; instead, we have to analyze the game in its original representation.

In this paper, we study how to compute the optimal strategy to commit to, both in normal-form games (Section 2) and in Bayesian games, which are a special case of extensive-form games (Section 3).

2. NORMAL-FORM GAMES

In this section, we study how to compute the optimal strategy to commit to for games represented in normal form.

2.1 Definitions

In a *normal-form game*, every player $i \in \{1, \dots, n\}$ has a set of *pure strategies* (or *actions*) S_i , and a utility function $u_i : S_1 \times S_2 \times \dots \times S_n \rightarrow \mathbb{R}$ that maps every *outcome* (a vector consisting of a pure strategy for every player, also known as a *profile* of pure strategies) to a real number. To ease notation, in the case of two players, we will refer to player 1’s pure strategy set as S , and player 2’s pure strategy set as T . Such games can be represented in (*bi-*)*matrix form*, in which the rows correspond to player 1’s pure strategies, the columns correspond to player 2’s pure strategies, and the entries of the matrix give the row and column player’s utilities (in that order) for the corresponding outcome of the game. In the case of three players, we will use R , S , and T , for player 1, 2, and 3’s pure strategies, respectively. A *mixed strategy* for a player is a probability distribution over that player’s pure strategies. In the case of two-player games, we will refer to player 1 as the *leader* and player 2 as the *follower*.

Before defining optimal leadership strategies, consider the following game which illustrates the effect of the leader’s ability to commit.

2, 1	4, 0
1, 0	3, 1

In this normal-form representation, the bottom strategy for the row player is strictly dominated by the top strategy. Nevertheless, if the row player has the ability to commit to a pure strategy before the column player chooses his strategy, the row player should commit to the bottom strategy: doing so will make the column player prefer to play the right strategy, leading to a utility of 3 for the row player. By contrast, if the row player were to commit to the top strategy, the column player would prefer to play the left strategy, leading to a utility of only 2 for the row player. If the row player is able to commit to a mixed strategy, then she can get an even greater (expected) utility: if the row player commits to placing probability $p > 1/2$ on the bottom strategy, then the column player will still prefer to play the right strategy, and the row player’s expected utility will be $3p + 4(1 - p) = 4 - p \geq 3$. If the row player plays each strategy with probability exactly $1/2$, the column player is

indifferent between the strategies. In such cases, we will assume that the column player will choose the strategy that maximizes the row player's utility (in this case, the right strategy). Hence, the optimal mixed strategy to commit to for the row player is $p = 1/2$. There are a few good reasons for this assumption. If we were to assume the opposite, then there would not exist an optimal strategy for the row player in the example game: the row player would play the bottom strategy with probability $p = 1/2 + \epsilon$ with $\epsilon > 0$, and the smaller ϵ , the better the utility for the row player. By contrast, if we assume that the follower always breaks ties in the leader's favor, then an optimal mixed strategy for the leader always exists, and this corresponds to a subgame perfect equilibrium of the extensive-form representation of the leadership situation. In any case, this is a standard assumption for such models (e.g. [20]), although some work has investigated what can happen in the other subgame perfect equilibria [26]. (For *generic* two-player games, the leader's subgame-perfect equilibrium payoff is unique.) Also, the same assumption is typically used in mechanism design, in that it is assumed that if an agent is indifferent between revealing his preferences truthfully and revealing them falsely, he will report them truthfully. Given this assumption, we can safely refer to "optimal leadership strategies" rather than having to use some equilibrium notion.

Hence, for the purposes of this paper, an *optimal strategy to commit to* in a 2-player game is a strategy $s \in S'$ that maximizes $\max_{t \in BR(s)} u_l(s, t)$, where $BR(s) = \arg \max_{t \in T} u_f(s, t)$. (u_l and u_f are the leader and follower's utility functions, respectively.) We can have $S' = S$ for the case of commitment to pure strategies, or $S' = \Delta(S)$, the set of probability distributions over S , for the case of commitment to mixed strategies. (We note that replacing T by $\Delta(T)$ makes no difference in this definition.) For games with more than two players, in which the players commit to their strategies in sequence, we define optimal strategies to commit to recursively. After the leader commits to a strategy, the game to be played by the remaining agents is itself a (smaller) leadership game. Thus, we define an optimal strategy to commit to as a strategy that maximizes the leader's utility, assuming that the play of the remaining agents is itself optimal under this definition, and maximizes the leader's utility among all optimal ways to play the remaining game. Again, commitment to mixed strategies may or may not be a possibility for every player (although for the last player it does not matter if we allow for commitment to mixed strategies).

2.2 Commitment to pure strategies

We first study how to compute the optimal *pure* strategy to commit to. This is relatively simple, because the number of strategies to commit to is not very large. (In the following, *#outcomes* is the number of complete strategy profiles.)

THEOREM 1. *Under commitment to pure strategies, the set of all optimal strategy profiles in a normal-form game can be found in $O(\#players \cdot \#outcomes)$ time.*

PROOF. Each pure strategy that the first player may commit to will induce a subgame for the remaining players. We can solve each such subgame recursively to find all of its optimal strategy profiles; each of these will give the original leader some utility. Those that give the leader maximal utility correspond exactly to the optimal strategy profiles of the original game.

We now present the algorithm formally. Let $Su(G, s_1)$ be the subgame that results after the first (remaining) player in G plays $s_1 \in S_1^G$. A game with 0 players is simply an outcome of the game. The function $\text{Append}(s, O)$ appends the strategy s to each of the vectors of strategies in the set O . Let e be the empty vector with no elements. In a slight abuse of notation, we will write $u_1^G(C)$ when all strategy profiles in the set C give player 1 the same utility in the game G . (Here, player 1 is the first remaining player in the subgame G , not necessarily player 1 in the original game.) We note that $\arg \max$ is set-valued. Then, the following algorithm computes all optimal strategy profiles:

```

Algorithm Solve( $G$ )
if  $G$  has 0 players
  return  $\{e\}$ 
 $C \leftarrow \emptyset$ 
for all  $s_1 \in S_1^G$  {
   $O \leftarrow \text{Solve}(Su(G, s_1))$ 
   $O' \leftarrow \arg \max_{o \in O} u_1^G(s_1, o)$ 
  if  $C = \emptyset$  or  $u_1^G(s_1, O') = u_1^G(C)$ 
     $C \leftarrow C \cup \text{Append}(s_1, O')$ 
  if  $u_1^G(s_1, O') > u_1^G(C)$ 
     $C \leftarrow \text{Append}(s_1, O')$ 
}
return  $C$ 

```

Every outcome is (potentially) examined by every player, which leads to the given runtime bound. \square

As an example of how the algorithm works, consider the following 3-player game, in which the first player chooses the left or right matrix, the second player chooses a row, and the third player chooses a column.

0,1,1	1,1,0	1,0,1	3,3,0	0,2,0	3,0,1
2,1,1	3,0,1	1,1,1	4,4,2	0,0,2	0,0,0
0,0,1	0,0,0	3,3,0	0,5,1	0,0,0	3,0,0

First we eliminate the outcomes that do not correspond to best responses for the third player (removing them from the matrix):

0,1,1		1,0,1			3,0,1
2,1,1	3,0,1	1,1,1	4,4,2	0,0,2	
0,0,1			0,5,1		

Next, we remove the entries in which the third player does not break ties in favor of the second player, as well as entries that do not correspond to best responses for the second player.

0,1,1					
2,1,1		1,1,1			
			0,5,1		

Finally, we remove the entries in which the second and third players do not break ties in favor of the first player, as well as entries that do not correspond to best responses for the first player.

2,1,1					

Hence, in optimal play, the first player chooses the left matrix, the second player chooses the middle row, and the third player chooses the left column. (We note that this outcome is Pareto-dominated by (Right, Middle, Left).)

For general normal-form games, each player’s utility for each of the outcomes has to be explicitly represented in the input, so that the input size is itself $\Omega(\#players \cdot \#outcomes)$. Therefore, the algorithm is in fact a linear-time algorithm.

2.3 Commitment to mixed strategies

In the special case of two-player zero-sum games, computing an optimal mixed strategy for the leader to commit to is equivalent to computing a *minimax* strategy, which minimizes the maximum expected utility that the opponent can obtain. Minimax strategies constitute the only natural solution concept for two-player zero-sum games: von Neumann’s Minimax Theorem [24] states that in two-player zero-sum games, it does not matter (in terms of the players’ utilities) which player gets to commit to a mixed strategy first, and a profile of mixed strategies is a Nash equilibrium if and only if both strategies are minimax strategies. It is well-known that a minimax strategy can be found in polynomial time, using linear programming [17]. Our first result in this section generalizes this result, showing that an optimal mixed strategy for the leader to commit to can be efficiently computed in general-sum two-player games, again using linear programming.

THEOREM 2. *In 2-player normal-form games, an optimal mixed strategy to commit to can be found in polynomial time using linear programming.*

PROOF. For every pure follower strategy t , we compute a mixed strategy for the leader such that 1) playing t is a best response for the follower, and 2) under this constraint, the mixed strategy maximizes the leader’s utility. Such a mixed strategy can be computed using the following simple linear program:

$$\begin{array}{l} \text{maximize } \sum_{s \in S} p_s u_l(s, t) \\ \text{subject to} \\ \text{for all } t' \in T, \sum_{s \in S} p_s u_f(s, t) \geq \sum_{s \in S} p_s u_f(s, t') \\ \sum_{s \in S} p_s = 1 \end{array}$$

We note that this program may be infeasible for some follower strategies t , for example, if t is a strictly dominated strategy. Nevertheless, the program must be feasible for at least some follower strategies; among these follower strategies, choose a strategy t^* that maximizes the linear program’s solution value. Then, if the leader chooses as her mixed strategy the optimal settings of the variables p_s for the linear program for t^* , and the follower plays t^* , this constitutes an optimal strategy profile. \square

In the following result, we show that we cannot expect to solve the problem more efficiently than linear programming, because we can reduce any linear program with a probability constraint on its variables to a problem of computing the optimal mixed strategy to commit to in a 2-player normal-form game.

THEOREM 3. *Any linear program whose variables x_i (with $x_i \in \mathbb{R}^{\geq 0}$) must satisfy $\sum_i x_i = 1$ can be modeled as a prob-*

lem of computing the optimal mixed strategy to commit to in a 2-player normal-form game.

PROOF. Let the leader have a pure strategy i for every variable x_i . Let the column player have one pure strategy j for every constraint in the linear program (other than $\sum_i x_i = 1$), and a single additional pure strategy 0. Let the utility functions be as follows. Writing the objective of the linear program as **maximize** $\sum_i c_i x_i$, for any i , let $u_l(i, 0) = c_i$ and $u_f(i, 0) = 0$. Writing the j th constraint of the linear program (not including $\sum_i x_i = 1$) as $\sum_i a_{ij} x_i \leq b_j$, for any $i, j > 0$, let $u_l(i, j) = \min_{i'} c_{i'} - 1$ and $u_f(i, j) = a_{ij} - b_j$.

For example, consider the following linear program.

$$\begin{array}{l} \text{maximize } 2x_1 + x_2 \\ \text{subject to} \\ x_1 + x_2 = 1 \\ 5x_1 + 2x_2 \leq 3 \\ 7x_1 - 2x_2 \leq 2 \end{array}$$

The optimal solution to this program is $x_1 = 1/3$, $x_2 = 2/3$. Our reduction transforms this program into the following leader-follower game (where the leader is the row player).

2, 0	0, 2	0, 5
1, 0	0, -1	0, -4

Indeed, the optimal strategy for the leader is to play the top strategy with probability $1/3$ and the bottom strategy with probability $2/3$. We now show that the reduction works in general.

Clearly, the leader wants to incentivize the follower to play 0, because the utility that the leader gets when the follower plays 0 is always greater than when the follower does not play 0. In order for the follower not to prefer playing $j > 0$ rather than 0, it must be the case that $\sum_i p_l(i)(a_{ij} - b_j) \leq 0$, or equivalently $\sum_i p_l(i)a_{ij} \leq b_j$. Hence the leader will get a utility of at least $\min_{i'} c_{i'}$ if and only if there is a feasible solution to the constraints. Given that the $p_l(i)$ incentivize the follower to play 0, the leader attempts to maximize $\sum_i p_l(i)c_i$. Thus the leader must solve the original linear program. \square

As an alternative proof of Theorem 3, one may observe that it is known that finding a minimax strategy in a zero-sum game is as hard as the linear programming problem [6], and as we pointed out at the beginning of this section, computing a minimax strategy in a zero-sum game is a special case of the problem of computing an optimal mixed strategy to commit to.

This polynomial-time solvability of the problem of computing an optimal mixed strategy to commit to in two-player normal-form games contrasts with the unknown complexity of computing a Nash equilibrium in such games [21], as well as with the NP-hardness of finding a Nash equilibrium with maximum utility for a given player in such games [8, 2].

Unfortunately, this result does not generalize to more than two players—here, the problem becomes NP-hard. To show this, we reduce from the VERTEX-COVER problem.

DEFINITION 1. *In VERTEX-COVER, we are given a graph $G = (V, E)$ and an integer K . We are asked whether there*

exists a subset of the vertices $S \subseteq V$, with $|S| = K$, such that every edge $e \in E$ has at least one of its endpoints in S . *BALANCED-VERTEX-COVER* is the special case of *VERTEX-COVER* in which $K = |V|/2$.

VERTEX-COVER is NP-complete [9]. The following lemma shows that the hardness remains if we require $K = |V|/2$. (Similar results have been shown for other NP-complete problems.)

LEMMA 1. *BALANCED-VERTEX-COVER* is NP-complete.

PROOF. Membership in NP follows from the fact that the problem is a special case of *VERTEX-COVER*, which is in NP. To show NP-hardness, we reduce an arbitrary *VERTEX-COVER* instance to a *BALANCED-VERTEX-COVER* instance, as follows. If, for the *VERTEX-COVER* instance, $K > |V|/2$, then we simply add isolated vertices that are disjoint from the rest of the graph, until $K = |V|/2$. If $K < |V|/2$, we add isolated triangles (that is, the complete graph on three vertices) to the graph, increasing K by 2 every time, until $K = |V|/2$. \square

THEOREM 4. In 3-player normal-form games, finding an optimal mixed strategy to commit to is NP-hard.

PROOF. We reduce an arbitrary *BALANCED-VERTEX-COVER* instance to the following 3-player normal-form game. For every vertex v , each of the three players has a pure strategy corresponding to that vertex (r_v, s_v, t_v , respectively). In addition, for every edge e , the third player has a pure strategy t_e ; and finally, the third player has one additional pure strategy t_0 . The utilities are as follows:

- for all $r \in R, s \in S, u_1(r, s, t_0) = u_2(r, s, t_0) = 1$;
- for all $r \in R, s \in S, t \in T - \{t_0\}, u_1(r, s, t) = u_2(r, s, t) = 0$;
- for all $v \in V, s \in S, u_3(r_v, s, t_v) = 0$;
- for all $v \in V, r \in R, u_3(r, s_v, t_v) = 0$;
- for all $v \in V$, for all $r \in R - \{r_v\}, s \in S - \{s_v\}, u_3(r, s, t_v) = \frac{|V|}{|V|-2}$;
- for all $e \in E, s \in S$, for both $v \in e, u_3(r_v, s, t_e) = 0$;
- for all $e \in E, s \in S$, for all $v \notin e, u_3(r_v, s, t_e) = \frac{|V|}{|V|-2}$.
- for all $r \in R, s \in S, u_3(r, s, t_0) = 1$.

We note that players 1 and 2 have the same utility function. We claim that there is an optimal strategy profile in which players 1 and 2 both obtain 1 (their maximum utility) if and only if there is a solution to the *BALANCED-VERTEX-COVER* problem. (Otherwise, these players will both obtain 0.)

First, suppose there exists a solution to the *BALANCED-VERTEX-COVER* problem. Then, let player 1 play every r_v such that v is in the cover with probability $\frac{2}{|V|}$, and let player 2 play every s_v such that v is *not* in the cover with probability $\frac{2}{|V|}$. Then, for player 3, the expected utility of playing t_v (for any v) is $(1 - \frac{2}{|V|})\frac{|V|}{|V|-2} = 1$, because there is a chance of $\frac{2}{|V|}$ that r_v or s_v is played. Additionally, the expected utility of playing t_e (for any e) is at most

$(1 - \frac{2}{|V|})\frac{|V|}{|V|-2} = 1$, because there is a chance of at least $\frac{2}{|V|}$ that some r_v with $v \in e$ is played (because player 1 is randomizing over the pure strategies corresponding to the cover). It follows that playing t_0 is a best response for player 3, giving players 1 and 2 a utility of 1.

Now, suppose that players 1 and 2 obtain 1 in optimal play. Then, it must be the case that player 3 plays t_0 . Hence, for every $v \in V$, there must be a probability of at least $\frac{2}{|V|}$ that either r_v or s_v is played, for otherwise player 3 would be better off playing t_v . Because players 1 and 2 have only a total probability of 2 to distribute, it must be the case that for each v , *either* r_v *or* s_v is played with probability $\frac{2}{|V|}$, and the other is played with probability 0. (It is not possible for both to have nonzero probability, because then there would be some probability that both are played simultaneously (correlation is not possible), hence the total probability of at least one being played could not be high enough for all vertices.) Thus, for exactly half the $v \in V$, player 1 places probability $\frac{2}{|V|}$ on r_v . Moreover, for every $e \in E$, there must be a probability of at least $\frac{2}{|V|}$ that some r_v with $v \in e$ is played, for otherwise player 3 would be better off playing t_e . Thus, the $v \in V$ such that player 1 places probability $\frac{2}{|V|}$ on r_v constitute a balanced vertex cover. \square

3. BAYESIAN GAMES

So far, we have restricted our attention to normal-form games. In a normal-form game, it is assumed that every agent knows every other agent's preferences over the outcomes of the game. In general, however, agents may have some *private information* about their preferences that is not known to the other agents. Moreover, at the time of commitment to a strategy, the agents may not even know their *own* (final) preferences over the outcomes of the game yet, because these preferences may be dependent on a context that has yet to materialize. For example, when the code for a trading agent is written, it may not yet be clear how that agent will value resources that it will negotiate over later, because this depends on information that is not yet available at the time at which the code is written (such as orders that will have been placed to the agent before the negotiation). In this section, we will study commitment in *Bayesian* games, which can model such uncertainty over preferences.

3.1 Definitions

In a *Bayesian game*, every player i has a set of actions S_i , a set of *types* Θ_i with an associated probability distribution $\pi_i : \Theta_i \rightarrow [0, 1]$, and, for each type θ_i , a utility function $u_i^{\theta_i} : S_1 \times S_2 \times \dots \times S_n \rightarrow \mathbb{R}$. A *pure strategy* in a Bayesian game is a mapping from the player's types to actions, $\sigma_i : \Theta_i \rightarrow S_i$. (Bayesian games can be rewritten in normal form by enumerating every pure strategy σ_i , but this will cause an exponential blowup in the size of the representation of the game and therefore cannot lead to efficient algorithms.)

The strategy that the leader should commit to depends on whether, at the time of commitment, the leader knows her own type. If the leader does know her own type, the other types that the leader might have had become irrelevant and the leader should simply commit to the strategy that is optimal for the type. However, as argued above, the leader does not necessarily know her own type at the time of commitment (*e.g.*, the time at which the code is submitted). In this case, the leader must commit to a strategy that is

dependent upon the leader’s eventual type. We will study this latter model, although we will pay specific attention to the case where the leader has only a single type, which is effectively the same as the former model.

3.2 Commitment to pure strategies

It turns out that computing an optimal pure strategy to commit to is hard in Bayesian games, even with two players.

THEOREM 5. *Finding an optimal pure strategy to commit to in 2-player Bayesian games is NP-hard, even when the follower has only a single type.*

PROOF. We reduce an arbitrary VERTEX-COVER instance to the following Bayesian game between the leader and the follower. The leader has K types $\theta_1, \theta_2, \dots, \theta_K$, each occurring with probability $1/K$, and for every vertex $v \in V$, the leader has an action s_v . The follower has only a single type; for each edge $e \in E$, the follower has an action t_e , and the follower has a single additional action t_0 . The utility function for the leader is given by, for all $\theta_l \in \Theta_l$ and all $s \in S$, $u_l^{\theta_l}(s, t_0) = 1$, and for all $e \in E$, $u_l^{\theta_l}(s, t_e) = 0$. The follower’s utility is given by:

- For all $v \in V$, for all $e \in E$ with $v \notin e$, $u_f(s_v, t_e) = 1$;
- For all $v \in V$, for all $e \in E$ with $v \in e$, $u_f(s_v, t_e) = -K$;
- For all $v \in V$, $u_f(s_v, t_0) = 0$.

We claim that the leader can get a utility of 1 if and only if there is a solution to the VERTEX-COVER instance.

First, suppose that there is a solution to the VERTEX-COVER instance. Then, the leader can commit to a pure strategy such that for each vertex v in the cover, the leader plays s_v for some type. Then, the follower’s utility for playing t_e (for any $e \in E$) is at most $\frac{K-1}{K} + \frac{1}{K}(-K) = -\frac{1}{K}$, so that the follower will prefer to play t_0 , which gives the leader a utility of 1, as required.

Now, suppose that there is a pure strategy for the leader that will give the leader a utility of 1. Then, the follower must play t_0 . In order for the follower not to prefer playing t_e (for any $e \in E$) instead, for at least one $v \in e$ the leader must play s_v for some type θ_l . Hence, the set of vertices v that the leader plays for some type must constitute a vertex cover; and this set can have size at most K , because the leader has only K types. So there is a solution to the VERTEX-COVER instance. \square

However, if the *leader* has only a single type, then the problem becomes easy again ($\#types$ is the number of types for the follower):

THEOREM 6. *In 2-player Bayesian games in which the leader has only a single type, an optimal pure strategy to commit to can be found in $O(\#\text{outcomes} \cdot \#\text{types})$ time.*

PROOF. For every leader action s , we can compute, for every follower type $\theta_f \in \Theta_f$, which actions t maximize the follower’s utility; call this set of actions $BR_{\theta_f}(s)$. Then, the utility that the leader receives for committing to action s can be computed as $\sum_{\theta_f \in \Theta_f} \pi(\theta_f) \max_{t \in BR_{\theta_f}(s)} u_l(s, t)$, and the leader can choose the best action to commit to. \square

3.3 Commitment to mixed strategies

In two-player zero-sum imperfect information games with *perfect recall* (no player ever forgets something that it once knew), a minimax strategy can be constructed in polynomial time [12, 13]. Unfortunately, this result does not extend to computing optimal mixed strategies to commit to in the general-sum case—not even in Bayesian games. We will exhibit NP-hardness by reducing from the INDEPENDENT-SET problem.

DEFINITION 2. *In INDEPENDENT-SET, we are given a graph $G = (V, E)$ and an integer K . We are asked whether there exists a subset of the vertices $S \subseteq V$, with $|S| = K$, such that no edge $e \in E$ has both of its endpoints in S .*

Again, this problem is NP-complete [9].

THEOREM 7. *Finding an optimal mixed strategy to commit to in 2-player Bayesian games is NP-hard, even when the leader has only a single type and the follower has only two actions.*

PROOF. We reduce an arbitrary INDEPENDENT-SET instance to the following Bayesian game between the leader and the follower. The leader has only a single type, and for every vertex $v \in V$, the leader has an action s_v . The follower has a type θ_v for every $v \in V$, occurring with probability $\frac{1}{(|E|+1)|V|}$, and a type θ_e for every $e \in E$, occurring with probability $\frac{1}{|E|+1}$. The follower has two actions: t_0 and t_1 . The leader’s utility is given by, for all $s \in S$, $u_l(s, t_0) = 1$ and $u_l(s, t_1) = 0$. The follower’s utility is given by:

- For all $v \in V$, $u_f^{\theta_v}(s_v, t_1) = 0$;
- For all $v \in V$ and $s \in S - \{s_v\}$, $u_f^{\theta_v}(s, t_1) = \frac{K}{K-1}$;
- For all $v \in V$ and $s \in S$, $u_f^{\theta_v}(s, t_0) = 1$;
- For all $e \in E$, $s \in S$, $u_f^{\theta_e}(s, t_0) = 1$;
- For all $e \in E$, for both $v \in e$, $u_f^{\theta_e}(s_v, t_1) = \frac{2K}{3}$;
- For all $e \in E$, for all $v \notin e$, $u_f^{\theta_e}(s_v, t_1) = 0$.

We claim that an optimal strategy to commit to gives the leader an expected utility of at least $\frac{|E|}{|E|+1} + \frac{K}{(|E|+1)|V|}$ if and only if there is a solution to the INDEPENDENT-SET instance.

First, suppose that there is a solution to the INDEPENDENT-SET instance. Then, the leader could commit to the following strategy: for every vertex v in the independent set, play the corresponding s_v with probability $1/K$. If the follower has type θ_e for some $e \in E$, the expected utility for the follower of playing t_1 is at most $\frac{1}{K} \frac{2K}{3} = 2/3$, because there is at most one vertex $v \in e$ such that s_v is played with nonzero probability. Hence, the follower will play t_0 and obtain a utility of 1. If the follower has type θ_v for some vertex v in the independent set, the expected utility for the follower of playing t_1 is $\frac{K-1}{K} \frac{K}{K-1} = 1$, because the leader plays s_v with probability $1/K$. It follows that the follower (who breaks ties to maximize the leader’s utility) will play t_0 , which also gives a utility of 1 and gives the leader a higher utility. Hence the leader’s expected utility for this strategy is at least $\frac{|E|}{|E|+1} + \frac{K}{(|E|+1)|V|}$, as required.

Now, suppose that there is a strategy that gives the leader an expected utility of at least $\frac{|E|}{|E|+1} + \frac{K}{(|E|+1)|V|}$. Then, this strategy must induce the follower to play t_0 whenever it has a type of the form θ_e (because otherwise, the utility could be at most $\frac{|E|-1}{|E|+1} + \frac{|V|}{(|E|+1)|V|} = \frac{|E|}{|E|+1} < \frac{|E|}{|E|+1} + \frac{K}{(|E|+1)|V|}$). Thus, it cannot be the case that for some edge $e = (v_1, v_2) \in E$, the probability that the leader plays one of s_{v_1} and s_{v_2} is at least $2/K$, because then the expected utility for the follower of playing t_1 when it has type θ_e would be at least $\frac{2}{K} \frac{2K}{3} = 4/3 > 1$. Moreover, the strategy must induce the follower to play t_0 for at least K types of the form θ_v . Inducing the follower to play t_0 when it has type θ_v can be done only by playing s_v with probability at least $1/K$, which will give the follower a utility of at most $\frac{K-1}{K} \frac{K}{K-1} = 1$ for playing t_1 . But then, the set of vertices v such that s_v is played with probability at least $1/K$ must constitute an independent set of size K (because if there were an edge e between two such vertices, it would induce the follower to play t_1 for type θ_e by the above). \square

By contrast, if the *follower* has only a single type, then we can generalize the linear programming approach for normal-form games:

THEOREM 8. *In 2-player Bayesian games in which the follower has only a single type, an optimal mixed strategy to commit to can be found in polynomial time using linear programming.*

PROOF. We generalize the approach in Theorem 2 as follows. For every pure follower strategy t , we compute a mixed strategy for the leader for every one of the leader's types such that 1) playing t is a best response for the follower, and 2) under this constraint, the mixed strategy maximizes the leader's *ex ante* expected utility. To do so, we generalize the linear program as follows:

$$\begin{array}{l}
 \text{maximize } \sum_{\theta_l \in \Theta_l} \pi(\theta_l) \sum_{s \in S} p_s^{\theta_l} u_l^{\theta_l}(s, t) \\
 \text{subject to} \\
 \text{for all } t' \in T, \quad \sum_{\theta_l \in \Theta_l} \pi(\theta_l) \sum_{s \in S} p_s^{\theta_l} u_f(s, t) \geq \\
 \sum_{\theta_l \in \Theta_l} \pi(\theta_l) \sum_{s \in S} p_s^{\theta_l} u_f(s, t') \\
 \text{for all } \theta_l \in \Theta_l, \quad \sum_{s \in S} p_s^{\theta_l} = 1
 \end{array}$$

As in Theorem 2, the solution for the linear program that maximizes the solution value is an optimal strategy to commit to. \square

This shows an interesting contrast between commitment to pure strategies and commitment to mixed strategies in Bayesian games: for pure strategies, the problem becomes easy if the leader has only a single type (but not if the follower has only a single type), whereas for mixed strategies, the problem becomes easy if the follower has only a single type (but not if the leader has only a single type).

4. CONCLUSIONS AND FUTURE RESEARCH

In multiagent systems, strategic settings are often analyzed under the assumption that the players choose their strategies simultaneously. This requires some equilibrium

notion (Nash equilibrium and its refinements), and often leads to the *equilibrium selection* problem: it is unclear to each individual player according to which equilibrium she should play. However, this model is not always realistic. In many settings, one player is able to *commit* to a strategy before the other player makes a decision. For example, one agent may arrive at the (real or virtual) site of the game before the other, or, in the specific case of software agents, the code for one agent may be completed and committed before that of another agent. Such models are synonymously referred to as *leadership*, *commitment*, or *Stackelberg* models, and optimal play in such models is often significantly different from optimal play in the model where strategies are selected simultaneously. Specifically, if commitment to mixed strategies is possible, then (optimal) commitment never hurts the leader, and often helps.

The recent surge in interest in *computing* game-theoretic solutions has so far ignored leadership models (with the exception of the interest in *mechanism design*, where the designer is implicitly in a leadership position). In this paper, we studied how to compute optimal strategies to commit to under both commitment to pure strategies and commitment to mixed strategies, in both normal-form and Bayesian games. For normal-form games, we showed that the optimal pure strategy to commit to can be found efficiently for any number of players. An optimal mixed strategy to commit to in a normal-form game can be found efficiently for two players using linear programming (and no more efficiently than that, in the sense that any linear program with a probability constraint can be encoded as such a problem). (This is a generalization of the polynomial-time computability of minimax strategies in normal-form games.) The problem becomes NP-hard for three (or more) players. In Bayesian games, the problem of finding an optimal pure strategy to commit to is NP-hard even in two-player games in which the follower has only a single type, although two-player games in which the leader has only a single type can be solved efficiently. The problem of finding an optimal mixed strategy to commit to in a Bayesian game is NP-hard even in two-player games in which the leader has only a single type, although two-player games in which the follower has only a single type can be solved efficiently using a generalization of the linear programming approach for normal-form games. The following two tables summarize these results.

	2 players	≥ 3 players
normal-form	$O(\#outcomes)$	$O(\#outcomes \cdot \#players)$
Bayesian, 1-type leader	$O(\#outcomes \cdot \#types)$	NP-hard
Bayesian, 1-type follower	NP-hard	NP-hard
Bayesian (general)	NP-hard	NP-hard

Results for commitment to pure strategies. (With more than 2 players, the "follower" is the last player to commit, the "leader" is the first.)

	2 players	≥ 3 players
normal-form	one LP-solve per follower action	NP-hard
Bayesian, 1-type leader	NP-hard	NP-hard
Bayesian, 1-type follower	one LP-solve per follower action	NP-hard
Bayesian (general)	NP-hard	NP-hard

Results for commitment to mixed strategies. (With more than 2 players, the “follower” is the last player to commit, the “leader” is the first.)

Future research can take a number of directions. First, we can empirically evaluate the techniques presented here on test suites such as GAMUT [19]. We can also study the computation of optimal strategies to commit to in other¹ concise representations of normal-form games—for example, in graphical games [10] or local-effect/action graph games [14, 1]. For the cases where computing an optimal strategy to commit to is NP-hard, we can also study the computation of *approximately* optimal strategies to commit to. While the correct definition of an approximately optimal strategy is in this setting may appear simple at first—it should be a strategy that, if the following players play optimally, performs almost as well as the optimal strategy in expectation—this definition becomes problematic when we consider that the other players may also be playing only approximately optimally. One may also study models in which multiple (but not all) players commit at the same time.

Another interesting direction to pursue is to see if computing optimal mixed strategies to commit to can help us in, or otherwise shed light on, computing Nash equilibria. Often, optimal mixed strategies to commit to are also Nash equilibrium strategies (for example, in two-player zero-sum games this is always true), although this is not always the case (for example, as we already pointed out, sometimes the optimal strategy to commit to is a strictly dominated strategy, which can never be a Nash equilibrium strategy).

5. REFERENCES

- [1] N. A. R. Bhat and K. Leyton-Brown. Computing Nash equilibria of action-graph games. In *Proceedings of the 20th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, Banff, Canada, 2004.
- [2] V. Conitzer and T. Sandholm. Complexity results about Nash equilibria. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 765–771, Acapulco, Mexico, 2003.
- [3] V. Conitzer and T. Sandholm. Complexity of (iterated) dominance. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 88–97, Vancouver, Canada, 2005.
- [4] V. Conitzer and T. Sandholm. A generalized strategy eliminability criterion and computational methods for applying it. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 483–488, Pittsburgh, PA, USA, 2005.
- [5] A. A. Cournot. *Recherches sur les principes mathématiques de la théorie des richesses (Researches into the Mathematical Principles of the Theory of Wealth)*. Hachette, Paris, 1838.
- [6] G. Dantzig. A proof of the equivalence of the programming problem and the game problem. In T. Koopmans, editor, *Activity Analysis of Production and Allocation*, pages 330–335. John Wiley & Sons, 1951.
- [7] I. Gilboa, E. Kalai, and E. Zemel. The complexity of eliminating dominated strategies. *Mathematics of Operation Research*, 18:553–565, 1993.
- [8] I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1:80–93, 1989.
- [9] R. Karp. Reducibility among combinatorial problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, NY, 1972.
- [10] M. Kearns, M. Littman, and S. Singh. Graphical models for game theory. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2001.
- [11] D. E. Knuth, C. H. Papadimitriou, and J. N. Tsitsiklis. A note on strategy elimination in bimatrix games. *Operations Research Letters*, 7(3):103–107, 1988.
- [12] D. Koller and N. Megiddo. The complexity of two-person zero-sum games in extensive form. *Games and Economic Behavior*, 4(4):528–552, Oct. 1992.
- [13] D. Koller, N. Megiddo, and B. von Stengel. Efficient computation of equilibria for extensive two-person games. *Games and Economic Behavior*, 14(2):247–259, 1996.
- [14] K. Leyton-Brown and M. Tennenholtz. Local-effect games. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, Acapulco, Mexico, 2003.
- [15] R. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 36–41, San Diego, CA, 2003.
- [16] M. Littman and P. Stone. A polynomial-time Nash equilibrium algorithm for repeated games. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 48–54, San Diego, CA, 2003.
- [17] R. D. Luce and H. Raiffa. *Games and Decisions*. John Wiley and Sons, New York, 1957. Dover republication 1989.
- [18] J. Nash. Equilibrium points in n-person games. *Proc. of the National Academy of Sciences*, 36:48–49, 1950.
- [19] E. Nudelman, J. Wortman, K. Leyton-Brown, and Y. Shoham. Run the GAMUT: A comprehensive approach to evaluating game-theoretic algorithms. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, New York, NY, USA, 2004.
- [20] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- [21] C. Papadimitriou. Algorithms, games and the Internet. In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*, pages 749–753, 2001.

¹Bayesian games are one potentially concise representation of normal-form games.

- [22] R. Porter, E. Nudelman, and Y. Shoham. Simple search methods for finding a Nash equilibrium. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 664–669, San Jose, CA, USA, 2004.
- [23] T. Sandholm, A. Gilpin, and V. Conitzer. Mixed-integer programming methods for finding Nash equilibria. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 495–501, Pittsburgh, PA, USA, 2005.
- [24] J. von Neumann. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100:295–320, 1927.
- [25] H. von Stackelberg. *Marktform und Gleichgewicht*. Springer, Vienna, 1934.
- [26] B. von Stengel and S. Zamir. Leadership with commitment to mixed strategies. CDAM Research Report LSE-CDAM-2004-01, London School of Economics, Feb. 2004.