

Commitment to Correlated Strategies

Vincent Conitzer and Dmytro Korzhyk

Department of Computer Science, Duke University
Durham, NC 27708, USA
{conitzer, dima}@cs.duke.edu

Abstract

The standard approach to computing an optimal mixed strategy to commit to is based on solving a set of linear programs, one for each of the follower's pure strategies. We show that these linear programs can be naturally merged into a single linear program; that this linear program can be interpreted as a formulation for the optimal *correlated* strategy to commit to, giving an easy proof of a result by von Stengel and Zamir that the leader's utility is at least the utility she gets in any correlated equilibrium of the simultaneous-move game; and that this linear program can be extended to compute optimal correlated strategies to commit to in games of three or more players. (Unlike in two-player games, in games of three or more players, the notions of optimal mixed and correlated strategies to commit to are truly distinct.) We give examples, and provide experimental results that indicate that for 50×50 games, this approach is usually significantly faster than the multiple-LPs approach.

Introduction

Game theory provides a mathematical framework for rational action in settings with multiple agents. As such, algorithms for computing game-theoretic solutions are of great interest to the multiagent systems community in AI.

It has long been well known in game theory that being able to commit to a course of action before the other player(s) move(s)—often referred to as a *Stackelberg* model (von Stackelberg 1934)—can bestow significant advantages. In recent years, the problem of *committing* an optimal strategy to commit to has started to receive a significant amount of attention, especially in the multiagent systems community. In the initial paper (Conitzer and Sandholm 2006), a number of variants were studied, including commitment to pure and to mixed strategies, in normal-form and in Bayesian games. There have been several other papers making progress on versions of the problem that concern standard game-theoretic representations, including Bayesian games (Paruchuri et al. 2008; Letchford, Conitzer, and Munagala 2009) and extensive-form games (Letchford and Conitzer 2010). Perhaps the biggest impulse to this line of research, though, is due to

the use of these techniques in several security applications that have recently been deployed in the real world, including the randomized placement of checkpoints and canine units at Los Angeles International airport (Pita et al. 2009) and the assignment of Federal Air Marshals to flights (Tsai et al. 2009). These developments have inspired work on computing optimal mixed strategies to commit to in a specific class of games called *security games* (Kiekintveld et al. 2009; Korzhyk, Conitzer, and Parr 2010).

In this paper, we focus on what is arguably the most basic problem under the general topic of computing optimal mixed strategies to commit to: given a game represented in normal form, compute an optimal mixed strategy for player 1 to commit to. The following game is commonly used as an example for this.

	<i>L</i>	<i>R</i>
<i>U</i>	(1,1)	(3,0)
<i>D</i>	(0,0)	(2,1)

Without commitment, this game is solvable by iterated strict dominance: *U* strictly dominates *D* for player 1; after removing *D*, *L* strictly dominates *R* for player 2. So the iterated strict dominance outcome (and hence the only equilibrium outcome) is (*U*, *L*), resulting in a utility of 1 for player 1. However, if player 1 can commit to a pure strategy before player 2 moves, then player 1 is better off committing to *D*, thereby incentivizing player 2 to play *R*, resulting in a utility of 2 for player 1. Even better for player 1 is to commit to a mixed strategy of $(.49U, .51D)$; this still incentivizes player 2 to play *R* and results in an expected utility of $.49 \cdot 3 + .51 \cdot 2 = 2.49$ for player 1. Of course, it is even better to commit to $(.499U, .501D)$, etc. In the limit case of $(.5U, .5D)$, player 2 becomes indifferent between *L* and *R*; to guarantee the existence of an optimal solution, it is generally assumed that player 2 breaks ties in player 1's favor, so that $(.5U, .5D)$ is the unique optimal mixed strategy for player 1 to commit to, resulting in an expected utility of 2.5 for her.

It is already known (Conitzer and Sandholm 2006; von Stengel and Zamir 2010) that, in a two-player normal-form game, the optimal mixed strategy to commit to can be found in polynomial time, by solving multiple linear programs. (We will describe this approach in detail soon.) Still, given

the problem’s fundamental nature and importance, it seems worthwhile to investigate it in more detail. Can we design other, possibly more efficient algorithms? Can we relate this problem to other computational problems in game theory?

With three or more players, it is not immediately clear how to define the optimal mixed strategy to commit to for player 1. The reason is that, after player 1’s commitment, the remaining players play a game among themselves, and we need to consider according to which solution concept they will play. For example, to be consistent with the tie-breaking assumption in the two-player case, we could assume that the remaining players will play according to the Nash equilibrium of the remaining game that is best for player 1. However, this optimization problem is already NP-hard by itself, and in fact inapproximable unless $P=NP$ (Gilboa and Zemel 1989; Conitzer and Sandholm 2008), so there is little hope that this approach will lead to an efficient algorithm. Is there another natural solution concept that allows for a more efficient solution with three or more players?

In this paper, we make progress on these questions as follows. First, we show how to formulate the problem in the two-player setting as a single linear program,¹ and prove the correctness of this formulation by relating it to the existing multiple-LPs formulation. We then show how this single LP can be interpreted as a formulation for finding the optimal *correlated* strategy to commit to, giving an easy proof of a known result by von Stengel and Zamir (2010) that the optimal mixed strategy to commit to results in a utility for the leader that is at least as good as what she would obtain in any correlated equilibrium. We then show how this formulation can be extended to compute an optimal correlated strategy to commit to with three or more players, and illustrate by example that this can result in a higher utility for player 1 both compared to the best mixed strategy to commit to as well as compared to the best correlated equilibrium for player 1. (Unlike in two-player games, in games of three or more players, the notions of optimal mixed and correlated strategies to commit to are truly distinct.) Finally, we present experiments that indicate that, for 50×50 games drawn from “most” distribution families, our formulation is significantly faster than the multiple-LPs approach. We also investigate how often the correlated strategy is a product distribution (so that correlation does not play a role); as expected, with two players we always have a product distribution, but with more players this depends extremely strongly on the distribution over games that is used.

Review: multiple LPs

We now describe the standard approach to computing an optimal mixed strategy to commit to with two players (Conitzer and Sandholm 2006; von Stengel and Zamir 2010). The idea is a very natu-

¹Earlier work has already produced formulations of ((pre-)Bayesian versions of) the problem that involve only a single optimization (Paruchuri et al. 2008, extended version of Letchford, Conitzer, and Munagala 2009). However, these existing formulations use integer variables (even when restricted to the case of a single follower type).

ral divide-and-conquer approach: because we can assume without loss of optimality that in the solution, player 2 will play a pure strategy, we can simply consider each pure strategy for player 2 in turn. Let player i ’s set of pure strategies be S_i . For each pure strategy $s_2 \in S_2$ for player 2, we solve for the optimal mixed strategy for player 1, *under the constraint that s_2 is a best response for player 2*.

Linear Program 1 (known).

$$\begin{aligned} & \max \sum_{s_1 \in S_1} p_{s_1} u_1(s_1, s_2) \\ & \text{subject to:} \\ & (\forall s'_2 \in S_2) \sum_{s_1 \in S_1} p_{s_1} u_2(s_1, s'_2) \leq \sum_{s_1 \in S_1} p_{s_1} u_2(s_1, s_2) \\ & \sum_{s_1 \in S_1} p_{s_1} = 1 \\ & (\forall s_1 \in S_1) p_{s_1} \geq 0 \end{aligned}$$

(The first constraint says that player 2 should not be better off playing s'_2 instead of s_2 .) There is one of these linear programs for every s_2 , and at least one of these must have a feasible solution. We choose one with the highest optimal solution value; an optimal solution to this linear program corresponds to an optimal mixed strategy to commit to. Because linear programs can be solved in polynomial time, this gives a polynomial-time algorithm for computing an optimal mixed strategy to commit to.

A single linear program

Instead of solving one separate linear program per pure strategy for player 2, we can also solve the following single linear program:

Linear Program 2.

$$\begin{aligned} & \max \sum_{s_1 \in S_1, s_2 \in S_2} p_{(s_1, s_2)} u_1(s_1, s_2) \\ & \text{subject to:} \\ & (\forall s_2, s'_2 \in S_2) \\ & \quad \sum_{s_1 \in S_1} p_{(s_1, s_2)} u_2(s_1, s'_2) \leq \sum_{s_1 \in S_1} p_{(s_1, s_2)} u_2(s_1, s_2) \\ & \quad \sum_{s_1 \in S_1, s_2 \in S_2} p_{(s_1, s_2)} = 1 \\ & (\forall s_1 \in S_1, s_2 \in S_2) p_{(s_1, s_2)} \geq 0 \end{aligned}$$

We now explain why this linear program gives the right answer. The constraint matrix for Linear Program 2 has blocks along the diagonal: for each $s_2 \in S_2$, there is a set of constraints (one constraint for every $s'_2 \in S_2$) whose only nonzero coefficients correspond to the variables $p_{(s_1, s_2)}$ (one variable for every $s_1 \in S_1$). The exception is the probability constraint which has nonzero coefficients for all variables. (Cf. Dantzig-Wolfe decomposition.) The following proposition will help us to understand the relationship to the multiple-LPs approach, and hence the correctness of Linear Program 2.

Proposition 1 *Linear Program 2 always has an optimal solution in which only a single block of variables takes nonzero values. That is, there exists an optimal solution for which there is some $s_2^* \in S_2$ such that for any $s_1 \in S_1, s_2 \in S_2$ where $s_2 \neq s_2^*, p_{(s_1, s_2)} = 0$.*

Proof. Suppose for the sake of contradiction that all optimal solutions require nonzero values for at least k blocks, where $k \geq 2$. For an optimal solution p with exactly k nonzero blocks, let $s_2, s'_2 \in S_2$, $s_2 \neq s'_2$ be such that $t_{s_2} = \sum_{s_1 \in S_1} p_{(s_1, s_2)} > 0$ and $t_{s'_2} = \sum_{s_1 \in S_1} p_{(s_1, s'_2)} > 0$. Let $v_{s_2} = \sum_{s_1 \in S_1} p_{(s_1, s_2)} u_1(s_1, s_2)$ and $v_{s'_2} = \sum_{s_1 \in S_1} p_{(s_1, s'_2)} u_1(s_1, s'_2)$. Without loss of generality, suppose that $v_{s_2}/t_{s_2} \geq v_{s'_2}/t_{s'_2}$. Then consider the following modified solution p' :

- for all $s_1 \in S_1$, $p'_{(s_1, s_2)} = \frac{t_{s_2} + t_{s'_2}}{t_{s_2}} p_{(s_1, s_2)}$;
- for all $s_1 \in S_1$, $p'_{(s_1, s'_2)} = 0$;
- for all $s'_2 \notin \{s_2, s'_2\}$, $p'_{(s_1, s'_2)} = p_{(s_1, s'_2)}$.

p' has $k - 1$ blocks with nonzero values; we will show that p' remains feasible and has at least the same objective value as p , and must therefore be optimal, so that we arrive at the desired contradiction.

To prove that p' is still feasible, we first notice that any of the constraints corresponding to the unchanged blocks (for $s'_2 \notin \{s_2, s'_2\}$) must still hold because none of the variables with nonzero coefficients in these constraints have changed value. The constraints for the block corresponding to s'_2 hold trivially because all the variables with nonzero coefficients are set to zero. The constraints for the block corresponding to s_2 still hold because all the variables with nonzero coefficients have been multiplied by the same constant $\frac{t_{s_2} + t_{s'_2}}{t_{s_2}}$. Finally, the probability constraint still holds because the total probability on the variables in the s_2 block is $\sum_{s_1 \in S_1} p'_{(s_1, s_2)} = \sum_{s_1 \in S_1} \frac{t_{s_2} + t_{s'_2}}{t_{s_2}} p_{(s_1, s_2)} = \frac{t_{s_2} + t_{s'_2}}{t_{s_2}} t_{s_2} = t_{s_2} + t_{s'_2}$, that is, we have simply shifted the probability mass from s'_2 to s_2 . (All the probabilities are also still nonnegative, because $\frac{t_{s_2} + t_{s'_2}}{t_{s_2}}$ is positive.)

To prove that p' is no worse than p , we note that the total objective value derived under p' from variables in the s_2 block of variables is $\sum_{s_1 \in S_1} p'_{(s_1, s_2)} u_1(s_1, s_2) = \sum_{s_1 \in S_1} \frac{t_{s_2} + t_{s'_2}}{t_{s_2}} p_{(s_1, s_2)} u_1(s_1, s_2) = \frac{t_{s_2} + t_{s'_2}}{t_{s_2}} v_{s_2} \geq v_{s_2} + v_{s'_2}$, where the inequality follows from $v_{s_2}/t_{s_2} \geq v_{s'_2}/t_{s'_2}$. On the other hand, the total objective value derived under p' from variables in the s'_2 block of variables is 0 because all these variables are set to zero. In contrast, under the solution p , the total objective value from these two blocks is $v_{s_2} + v_{s'_2}$. Because p and p' agree on the other blocks, it follows that p' obtains at least as large an objective value as p , and we have a contradiction. \square

Proposition 1 suggests that one approach to solving Linear Program 2 is to force all the variables to zero with the exception of a single block and solve the remaining linear program; we try this for every block, and take the optimal solution overall. However, this approach coincides exactly with the original multiple-LPs approach, because:

Observation 1 In Linear Program 2, if for some s_2 , we force all the variables $p_{(s_1, s'_2)}$ for which $s'_2 \neq s_2$ to zero,

then the linear program that remains is identical to Linear Program 1.

This also proves the correctness of Linear Program 2 (because the multiple-LPs approach is correct).

Game-theoretic interpretation: commitment to correlated strategies

Linear Program 2 can be interpreted as follows. Player 1 commits to a *correlated* strategy. This entails that player 1 chooses a distribution $p_{(s_1, s_2)}$ over the outcomes, and commits to acting as follows: she draws (s_1, s_2) according to the distribution, recommends to player 2 that he should play s_2 , and plays s_1 herself. The constraints

$(\forall s_2, s'_2 \in S_2)$
 $\sum_{s_1 \in S_1} p_{(s_1, s_2)} u_2(s_1, s'_2) \leq \sum_{s_1 \in S_1} p_{(s_1, s_2)} u_2(s_1, s_2)$
in Linear Program 2 then mean that player 2 should always follow the recommendation s_2 rather than take some alternative action s'_2 . This is for the following reasons: if for some s_2 , $\sum_{s_1 \in S_1} p_{(s_1, s_2)} = 0$, then there will never be a recommendation to player 2 to play s_2 , and indeed the constraint will hold trivially in this case. On the other hand, if $\sum_{s_1 \in S_1} p_{(s_1, s_2)} > 0$, then player 2's subjective probability that player 1 will play s_1 given a recommendation of s_2 is $P(s_1|s_2) = \frac{p_{(s_1, s_2)}}{\sum_{s'_1 \in S_1} p_{(s'_1, s_2)}}$. Hence player 2 will be incentivized to follow the recommendation of playing s_2 rather than s'_2 if and only if

$(\forall s_2, s'_2 \in S_2)$
 $\sum_{s_1 \in S_1} P(s_1|s_2) u_2(s_1, s'_2) \leq \sum_{s_1 \in S_1} P(s_1|s_2) u_2(s_1, s_2)$
which is identical to the constraint in Linear Program 2 (by multiplying by $\sum_{s'_1 \in S_1} p_{(s'_1, s_2)}$).

Proposition 1 entails that we can without loss of optimality restrict attention to solutions where the recommendation to player 2 is always the same (so that there is no information in the signal to player 2).

Given this interpretation of Linear Program 2, it is not surprising that it is very similar to the linear feasibility formulation of the correlated equilibrium problem for two players. (In a correlated equilibrium (Aumann 1974), a third party known as a mediator draws (s_1, s_2) , recommends to each player i to play s_i without telling what the recommendation to the other player is, and it is optimal for each player to follow the recommendation.) The linear feasibility formulation of the correlated equilibrium problem is as follows (and the justification for the incentive constraints is the same as above):

Linear Program 3 (known).

(no objective required) subject to:

$$\begin{aligned}
& (\forall s_1, s'_1 \in S_1) \\
& \sum_{s_2 \in S_2} p_{(s_1, s_2)} u_1(s'_1, s_2) \leq \sum_{s_2 \in S_2} p_{(s_1, s_2)} u_1(s_1, s_2) \\
& (\forall s_2, s'_2 \in S_2) \\
& \sum_{s_1 \in S_1} p_{(s_1, s_2)} u_2(s_1, s'_2) \leq \sum_{s_1 \in S_1} p_{(s_1, s_2)} u_2(s_1, s_2) \\
& \sum_{s_1 \in S_1, s_2 \in S_2} p_{(s_1, s_2)} = 1 \\
& (\forall s_1 \in S_1, s_2 \in S_2) p_{(s_1, s_2)} \geq 0
\end{aligned}$$

Linear Program 2 is identical to Linear Program 3, except that in Linear Program 2 we have dropped the incentive constraints for player 1, and added an objective of maximizing player 1's expected utility. If we add the incentive constraints for player 1 back in, then we obtain a linear program for finding the correlated equilibrium that maximizes player 1's utility. Because adding constraints cannot increase the objective value of a maximization problem, we immediately obtain the following corollary:

Corollary 1 ((von Stengel and Zamir 2010)) *Player 1's expected utility from optimally committing to a mixed strategy is at least as high as her utility in any correlated equilibrium of the simultaneous-move game.*

Commitment to correlated strategies with more players

We have already seen that committing to a correlated strategy in a two-player game is in some sense not particularly interesting, because without loss of optimality player 2 will always get the same recommendation from player 1. However, the same is not true for games with $n \geq 3$ players, where player 1 commits to a correlated strategy and sends recommendations to players $2, \dots, n$, who then play simultaneously. We can easily extend Linear Program 2 to this case of n players (just as it is well known that Linear Program 3 can be extended to the case of n players):

Linear program 4.

$$\begin{aligned} & \max \sum_{s_1 \in S_1, \dots, s_n \in S_n} p(s_1, \dots, s_n) u_1(s_1, \dots, s_n) \\ & \text{subject to:} \\ & (\forall i \in \{2, \dots, n\}) (\forall s_i, s'_i \in S_i) \\ & \quad \sum_{s_{-i} \in S_{-i}} p(s_i, s_{-i}) u_i(s'_i, s_{-i}) \leq \sum_{s_{-i} \in S_{-i}} p(s_i, s_{-i}) u_i(s_i, s_{-i}) \\ & \quad \sum_{s_1 \in S_1, \dots, s_n \in S_n} p(s_1, \dots, s_n) = 1 \\ & (\forall s_1 \in S_1, \dots, s_n \in S_n) p(s_1, \dots, s_n) \geq 0 \end{aligned}$$

(Here, we followed the standard game theory notation of using $-i$ to refer to "the players other than i .") Again, Linear Program 4 is simply the standard linear feasibility program for correlated equilibrium, with the constraints for player 1 omitted and with an objective of maximizing player 1's expected utility. This immediately implies the following proposition:

Proposition 2 *The optimal correlated strategy to commit to in an n -player normal-form game can be found in time polynomial in the size of the input.*

The following example illustrates that if there are three or more players, then commitment to a correlated strategy can be strictly better for player 1 than commitment to a mixed strategy (as well as any correlated equilibrium of the simultaneous-move version of the game).

Example. Consider a three-player game between a wildlife refuge Manager (player 1, aka. M), a Lion (player 2, aka. L), and a wildlife Photographer (player 3, aka. P). There are four

locations in the game: A and B (two locations in the refuge that are out in the open), C (a safe hiding place for the lion), and D (the wildlife photographer's home). Each player must choose a location: M can choose between A and B , L between A , B , and C , and P between A , B , and D .

M wants L to come out into the open, and would prefer even more to be in the same place as L in order to study him. Specifically, M gets utility 2 if she is in the same location as L, 1 if L is at A or B but not at the same location as M, and 0 otherwise. L just wants to avoid contact with humans. Specifically, L gets utility 1 unless he is in the same location as another player, in which case he gets 0. P wants to get a close-up shot of L, but would rather stay home than go out and be unsuccessful. Specifically, P gets utility 2 for being in the same location as L, and otherwise 1 for being at D , and 0 otherwise.

A correlated strategy specifies a probability for every outcome, that is, every feasible triplet of locations for the players. We will show that the unique optimal correlated strategy for M to commit to is: $p_{(A,B,D)} = 1/2$, $p_{(B,A,D)} = 1/2$. That is, she flips a coin to determine whether to go to A or B , signals to L to go to the other location of the two, and always signals to P to stay home at D . This is not a correlated equilibrium of the simultaneous-move game, because M would be better off going to the same location as L. This does not pose a problem because M is *committing* to the strategy. The other two players are best-responding by following the recommendations: L is getting his maximum possible utility of 1; P (whose signal is always D and thus carries no information) is getting 1 for staying home, and switching to either A or B would still leave her with an expected utility of 1.

To see that M cannot do better, note that L can guarantee himself a utility of 1 by always choosing C , so there is no feasible solution where L has positive probability of being in the same location as another player. Hence, in any feasible solution, any outcome where M gets utility 2 has zero probability. All that remains to show is that there is no other feasible solution in which M always gets utility 1. In any such solution, L must always choose A or B . Also, in any feasible solution, P cannot play A or B with positive probability, because she can never be in the same location as L; hence, if she played A or B with positive probability, she would end up with an expected utility strictly below 1, whereas she can guarantee herself 1 just by choosing D . Because M also cannot be in the same location as L with positive probability, it follows that only $p_{(A,B,D)}$ and $p_{(B,A,D)}$ can be set to positive values. If one of them is set to a value greater than $1/2$, then P would be better off choosing the location where L is more than half the time. It follows that $p_{(A,B,D)} = 1/2$, $p_{(B,A,D)} = 1/2$ is the unique optimal solution. \square

Experiments

While Linear Programs 2 and 4 are arguably valuable from the viewpoint of improving our conceptual understanding of commitment, it is also worthwhile to investigate them as an algorithmic contribution. Linear Program 4 allows us

to do something we could not do before, namely, to compute an optimal correlated strategy to commit to in games with more than two players. This cannot be said about the special case of Linear Program 2, because we already had the multiple-LPs approach. But how does Linear Program 2 compare to the multiple-LPs approach? At least, it provides a slight implementation advantage, in the sense of not having to write code to iterate through multiple LPs. More interestingly, what is the effect on runtime of using it rather than the multiple-LPs approach? Of course, this depends on the LP solver. Does the solver benefit from having the problem decomposed into multiple LPs? Or does it benefit from seeing the whole problem at once?

To answer these questions, we evaluated our approach on GAMUT (Nudelman et al. 2004), which generates games according to a variety of distributions. Focusing on two-player games, we used CPLEX 10.010 both for the multiple-LPs approach (LP1) and for Linear Program 2 (LP2). For each GAMUT game class, we generated 50 two-player games with 50 strategies per player and compared the time it takes to find the optimal strategies to commit to in these games using LP1 and LP2. We show the boxplots of the run times in Figure 1. Perhaps surprisingly, it turns out that LP2 generally solves much faster. One possible explanation for this is as follows. In the multiple-LPs approach, each block is solved to optimality separately. In contrast, when presented with LP2, the solver sees the entire problem instance all at once, which in principle could allow it to quickly prune some blocks as being clearly suboptimal. The only distribution for which LP2 is slower is RandomZeroSum. Unfortunately, preliminary experiments on random games indicate that LP2 does not scale gracefully to larger games, and that perhaps LP1 scales a little better. We conjecture that this is due to heavier memory requirements for LP2.

Table 1 shows how often the correlated strategy to commit to computed by Linear Program 4 is a product distribution. We say that a distribution $p(s_1, \dots, s_n)$ is a product distribution iff it satisfies the following condition.

$$\begin{aligned} \forall i \in \{1, \dots, n\} \\ \forall s_i \in \{s_i'' \in S_i : p(s_i'') > 0\} \\ \forall s_i' \in \{s_i' \in S_i : p(s_i') > 0\} \\ \forall s_{-i} \in S_{-i} : p(s_{-i}|s_i) = p(s_{-i}|s_i') \end{aligned}$$

Low percentages here indicate that correlation plays a significant role. To compute this data, we generated 50 payoff matrices with 10 strategies per player for each combination of a GAMUT class and a number of players. In other words, if the leader commits to a product distribution over the strategy profiles, then the recommendation each of the followers gets from the correlated strategy does not give out any information about the recommendations that the other players receive. In games with two players, the correlated strategy computed by LP2 is always a product distribution, as expected by Proposition 1. For games with more than two players, in some distributions correlation does not play a big role, and in others it does.

We say that a correlated strategy is a degenerate distribution if its support size is 1. A degenerate distribution is a

Game class \ # players	2		3		4	
	P	D	P	D	P	D
BidirectionalLEG	1	.96	.9	.86	.84	.84
CovariantGame	1	.48	.64	.6	.68	.68
DispersionGame	1	1	1	1	1	1
GuessTwoThirdsAve	1	1	0	0	0	0
MajorityVoting	1	.88	1	1	1	1
MinimumEffortGame	1	1	1	1	1	1
RandomGame	1	.42	.16	.08	.02	.02
RandomGraphicalGame	1	.4	.22	.1	.02	.02
RandomLEG	1	1	.92	.92	.02	.02
TravelersDilemma	1	0	1	1	.02	.02
UniformLEG	1	.96	.88	.86	.02	.02

Table 1: For each game class and number of players, the two numbers shown are the fractions of product distributions (P) and the fractions of degenerate distributions (D) among the correlated strategies computed by LP4.

special case of a product distribution. As we can see from Table 1, a large fraction of computed product distributions are actually degenerate.

Conclusion

In this paper, we have shown that in two-player games, optimal mixed strategies to commit to can be computed using a single linear program that, experimentally, appears to solve 50×50 games drawn from most common distributions significantly faster than the known multiple-LPs approach. This single linear program has a close relationship to the linear program formulation for correlated equilibrium, allowing us to immediately obtain a known result about the relationship to correlated equilibrium as a corollary, and to extend the program to compute an optimal correlated strategy to compute for more than two players (and with more than two players, we indeed get nontrivial correlations).

The observation that a player’s ability to commit can be modeled simply by dropping the incentive constraints for that player in an equilibrium formulation, and adding an objective, can also be used for variants of the problem. For example, in an n -player game in which a given coalition of k players is able to make a joint commitment, all that is needed is to drop the incentive constraints for those k players and formulate an appropriate objective (such as the sum of those players’ utilities).

A natural direction for future research is to try to extend the methodology here to game representations other than the normal form. Significant results on the efficient computation of correlated equilibria in succinctly represented games have been obtained, though optimizing over the space of correlated equilibria (which is close to what we do in this paper) poses more challenges (Papadimitriou and Roughgarden 2008; Jiang and Leyton-Brown 2010).

Another direction for future research is to investigate more thoroughly scaling to larger games and address the memory requirements of LP2. It may be possible to achieve (some of) the benefits of LP2 that we have observed here in the multiple-LPs approach by giving the solver bounds from earlier LPs to quickly prune later, suboptimal LPs.

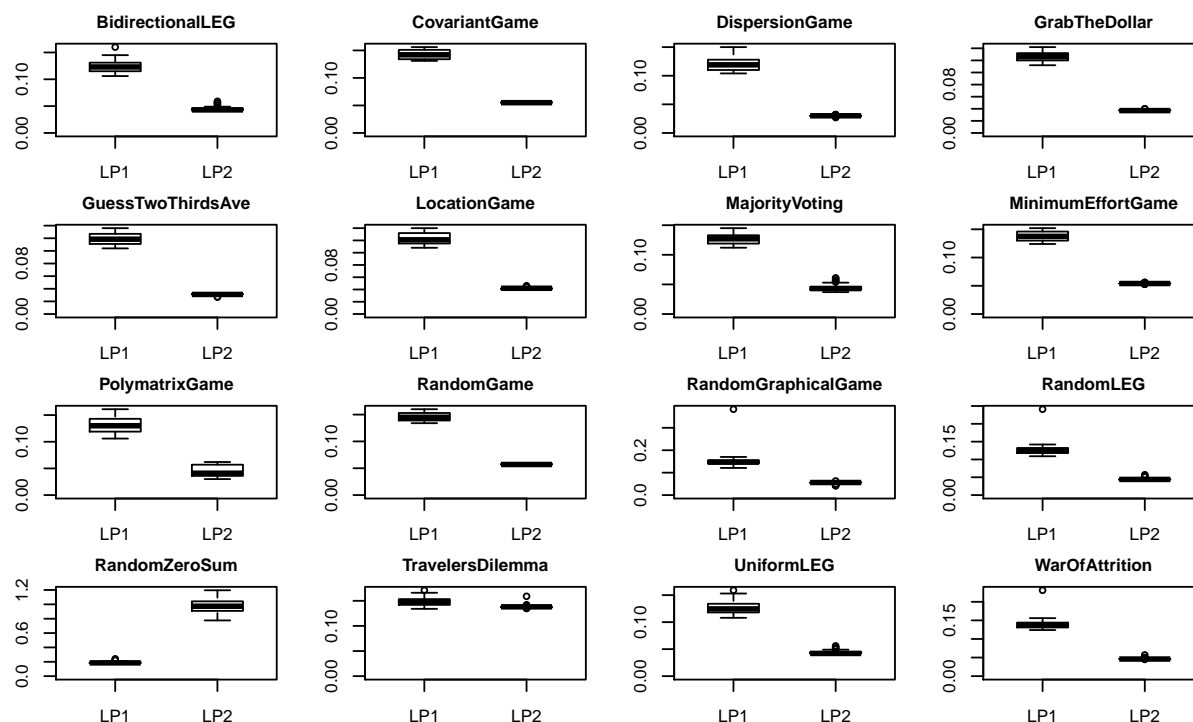


Figure 1: Run time comparison of LP1 and LP2 on GAMUT games (seconds).

Acknowledgments

We acknowledge NSF IIS-0812113, CAREER-0953756, ARO 56698-CI, and an Alfred P. Sloan Research Fellowship for support. We also thank Ronald Parr and the anonymous reviewers for their helpful comments.

References

- Aumann, R. 1974. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1:67–96.
- Conitzer, V., and Sandholm, T. 2006. Computing the optimal strategy to commit to. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, 82–90.
- Conitzer, V., and Sandholm, T. 2008. New complexity results about Nash equilibria. *Games and Economic Behavior* 63(2):621–641.
- Gilboa, I., and Zemel, E. 1989. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior* 1:80–93.
- Jiang, A. X., and Leyton-Brown, K. 2010. Polynomial-time computation of exact correlated equilibrium in compact games. arXiv:1011.0253v1.
- Kiekintveld, C.; Jain, M.; Tsai, J.; Pita, J.; Ordóñez, F.; and Tambe, M. 2009. Computing optimal randomized resource allocations for massive security games. In *Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 689–696.
- Korzhyk, D.; Conitzer, V.; and Parr, R. 2010. Complexity of computing optimal Stackelberg strategies in security resource allocation games. In *AAAI*, 805–810.
- Letchford, J., and Conitzer, V. 2010. Computing optimal strategies to commit to in extensive-form games. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, 83–92.
- Letchford, J.; Conitzer, V.; and Munagala, K. 2009. Learning and approximating the optimal strategy to commit to. In *Proceedings of the Second Symposium on Algorithmic Game Theory (SAGT-09)*, 250–262.
- Nudelman, E.; Wortman, J.; Leyton-Brown, K.; and Shoham, Y. 2004. Run the GAMUT: A comprehensive approach to evaluating game-theoretic algorithms. In *AAMAS*, 880–887.
- Papadimitriou, C. H., and Roughgarden, T. 2008. Computing correlated equilibria in multi-player games. *Journal of the ACM* 55(3).
- Paruchuri, P.; Pearce, J. P.; Marecki, J.; Tambe, M.; Ordóñez, F.; and Kraus, S. 2008. Playing games for security: An efficient exact algorithm for solving Bayesian Stackelberg games. In *AAMAS*, 895–902.
- Pita, J.; Jain, M.; Ordóñez, F.; Portway, C.; Tambe, M.; and Western, C. 2009. Using game theory for Los Angeles airport security. *AI Magazine* 30(1):43–57.
- Tsai, J.; Rathi, S.; Kiekintveld, C.; Ordóñez, F.; and Tambe, M. 2009. IRIS - a tool for strategic security allocation in transportation networks. In *AAMAS - Industry Track*, 37–44.
- von Stackelberg, H. 1934. *Marktform und Gleichgewicht*. Springer, Vienna. 58–70.
- von Stengel, B., and Zamir, S. 2010. Leadership games with convex strategy sets. *Games and Economic Behavior* 69:446–457.