Noname manuscript No.
(will be inserted by the editor)

On the Value of Commitment

Joshua Letchford · **Dmytro Korzhyk** · **Vincent Conitzer**

Received: date / Accepted: date

Abstract In game theory, it is well known that being able to commit to a strategy before other players move can be beneficial. In this paper, we analyze *how much* benefit a player can derive from commitment in various types of games, in a quantitative sense that is similar to concepts such as the value of mediation and the price of anarchy. Specifically, we introduce and study the *value of pure commitment* (the benefit of committing to a pure strategy), the *value of mixed commitment* (the benefit of committing to a mixed strategy), and the *mixed vs. pure commitment* ratio (how much can be gained by committing to a mixed strategy rather than a pure one). In addition to theoretical results about how large these values are in the extreme case in various classes of games, we also give average-case results based on randomly drawn normal-form games.

Keywords noncooperative game theory · commitment · Stackelberg · price of anarchy

J. Letchford
Duke University, Durham, NC, USA
Tel.: 919-660-6501
E-mail: jcl@cs.duke.edu

D. Korzhyk
Duke University, Durham, NC, USA
Tel.: 919-660-6501
E-mail: dima@cs.duke.edu

V. Conitzer
Duke University, Durham, NC, USA
Tel.: 919-660-6503
E-mail: conitzer@cs.duke.edu

1 Introduction

In settings where multiple self-interested agents (or *players*) interact in the same domain, game theory gives various solution concepts that prescribe how the players should act, such as Nash equilibrium, correlated equilibrium, etc. A *game* is defined as a set of players, the information and actions available to each of these players and a set of payoffs for each possible outcome. Sometimes, it is in the power of one of the players to effectively change the structure of the game. A notable example of this is that one player may be able to *commit* to a strategy (and communicate this commitment to the other players) before the other players move. While it may initially seem unintuitive that tying one's hands before the other players move could be beneficial, it is well known in game theory that this is indeed so. To see how this can affect the outcome of a game, consider the following game (which is often used as an example of this).

	L	R
U	(1,1)	(3,0)
D	(0,0)	(2,1)

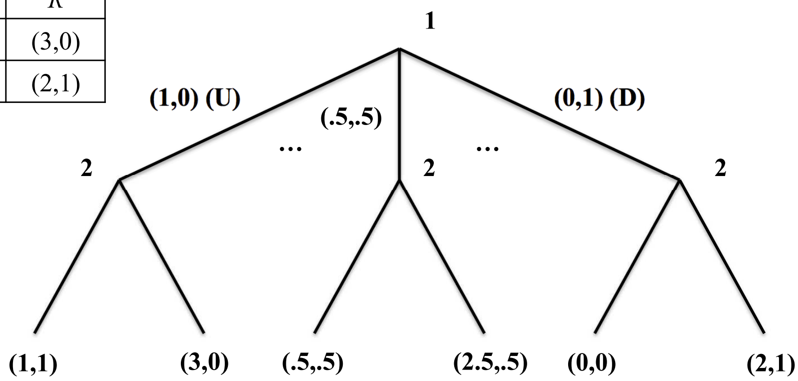

Fig. 1 A normal-form game, and the extensive-form representation of its commitment version.

Example 1 (known) Consider the normal-form game in Figure 1. For the case where the players move simultaneously (no ability to commit), the unique equilibrium is (U,L): U strictly dominates D, so that the game is solvable by iterated strict dominance. So, player 1 (the row player) receives utility 1. However, now suppose that player 1 has the ability to commit. Then, she is better off committing to play D, which will incentivize player 2 to play R, resulting in a utility of 2 for player 1. The situation gets even better for player 1 if she can commit to a mixed strategy: in this case, she can commit to the mixed strategy $(.5-\varepsilon, .5+\varepsilon)$, which still incentivizes player 2 to play R, but now player 1 receives an expected utility of $2.5-\varepsilon$. (Note that there is never a reason for player 2 to randomize, since he effectively faces a single-agent decision problem.) The game where player 1 can commit to a mixed strategy can be represented as an extensive-form game with a continuum of actions for player 1, as illustrated in Figure 1.

It has been previously shown that, in a sense, optimal commitment to a *mixed* strategy never hurts a player (relative to any Nash equilibrium or even any correlated equilibrium) [29]. (A rough intuition is that a player can simply commit to her equilibrium strategy.) In contrast, committing to a pure strategy is not always beneficial; for example, consider matching pennies.

An optimal strategy to commit to is usually called a *Stackelberg* strategy, after von Stackelberg, who showed that in Cournot's duopoly model [7], a firm that can commit to a production quantity has a strategic advantage [28]. The *computation* of Stackelberg strategies has recently started to receive significant attention. The first paper on this topic [6] appeared in EC-2006 and studied Stackelberg strategies in normal-form and Bayesian games, and showed, among other things, that the optimal mixed strategy to commit to in a two-player normal-form game can be found in polynomial time using linear programming,[1] though this becomes NP-hard in Bayesian games or with three players (a later paper [16] proved inapproximability results for the Bayesian case). Unfazed by the NP-hardness result, Paruchuri et al. [20] developed a mixed-integer programming formulation for the Bayesian case, and this algorithm has been implemented at the core of deployed security applications, specifically the strategic random placement of checkpoints and canine units at Los Angeles International Airport [9,21]. Later work [10,12] studied computing Stackelberg strategies in settings where the normal form has exponential size, for example, when player 1 has to allocate multiple resources to defend multiple targets. Additional motivating domains for this line of work are: the scheduling of Federal Air Marshals [27], Coast Guard patrols [26] and airport security on a national scale [22]. Additional work investigates computing optimal strategies to commit to in extensive-form [15] and stochastic games [17].

Real-world applications to security and law enforcement are not the only motivation for work on studying the effects of commitment; the notion of commitment plays a key role in many game-theoretic settings. Notably, in mechanism design (or environment design or principal-agent settings), the designer (or center, auctioneer, principal) is assumed to be able to commit to a mechanism before the (other) agents move. For example, if an auctioneer uses a Vickrey auction, it is generally assumed that she will unambiguously commit to this mechanism, rather than (for example) waiting for the bids to come in, and subsequently backtracking on her promise of a Vickrey auction and attempting to charge the winner her own bid (the first price) after all. As a result, the notion of commitment is a key link in the strong connection between game theory and mechanism design.

Our framework. In the literature discussed above, an advantage to committing to a strategy is often shown; moreover, we know that if commitment to a *mixed* strategy is possible, then at least in two-player games, committing to a mixed strategy never hurts [29]. However, to our knowledge, there has been no analysis of *how much* commitment can benefit a player. In this paper, we define and study the "value of commitment," which is the ratio between player 1's utility when she commits, and her utility when she does not commit. In the normal form game in Example 1, we considered two possibilities: player 1 can commit to a pure strategy, or to a mixed strategy. Without commitment, player 1 gets 1. With pure-strategy commitment, player 1 gets 2. Hence, the value of pure commitment (*VoPC*) for this game is $2/1 = 2$. With mixed-strategy commitment, player 1 gets 2.5. Hence, the value of mixed commitment (*VoMC*) for this game is $2.5/1 = 2.5$. We will also be interested in how much player 1 gains by committing to a mixed strategy rather than a pure one—this mixed vs. pure (*MvP*) ratio is $2.5/2 = 1.25$ for this game. In fact, we will generally be interested in *classes* of games, and the *largest* values that these ratios can attain. We will define these concepts formally in Section 2.

Related concepts and research. Perhaps the most closely related paper is "On the Value of Correlation" by Ashlagi et al. [2]. In this work, the authors consider the phenomenon that in some games, there are correlated equilibria that lead to higher welfare than any Nash equilibrium, and they study the ratio between the welfare in the best correlated equilibrium and the best Nash equilibrium. They call this ratio the *value of mediation*. This is similar to our notion of value of commitment, except that we consider player 1 (the leader)'s utility rather than the combined welfare of all the players, and are interested in the ratio between the solution with and without commitment (according to several solution concepts) rather than between correlated and Nash equilibrium. (Ashlagi et al. also study the ratio between the maximum possible welfare and the maximum welfare obtained in a correlated equilibrium.)

The work by Ashlagi et al., in turn, builds upon existing work that tries to evaluate the welfare obtained under certain solution concepts using certain ratios. Most notably, there is the work on the *price of anarchy* [14,19]. The price of anarchy is a measure of the amount of welfare lost due to the selfishness

[1] As is common in this literature, the algorithm presented in this paper assumes that ties are broken in the leader's favor. The same algorithm appears in a GEB-2010 paper by von Stengel and Zamir [29] to compute the highest possible payoff for the leader; they also give a linear-programming algorithm to compute the lowest possible payoff for the leader, and point out that these payoffs must be the same in generic bimatrix games.

of the players (relative to what a central planner could obtain). Specifically, it looks at the worst possible ratio between the worst Nash equilibrium and the optimal social welfare obtainable from the game. Of particular interest is the cost of selfish routing [24], which considers the ratio between the latency in Nash equilibrium and the optimal latency. There is also a literature on *Stackelberg routing* [23, 11]. While the use of the word "Stackelberg" here indicates a superficial similarity to our work (in particular the routing games that we study later in the paper), there is a fundamental difference. In the existing Stackelberg routing literature, a benevolent central authority first routes some of the flow before the selfish players move, and the goal of the central authority is to minimize the resulting price of anarchy (i.e., the ratio between the cost of the resulting flow and the optimal flow). In contrast, in our setting, player 1 is not a benevolent central authority, but rather a selfish player like any other, who is trying to minimize her own cost (or maximize her own utility), as is usually the case in Stackelberg games; and we are interested in the improvement in her own cost/utility that the player can obtain through commitment. One exception showed that for single commodity markets with nonatomic players and a Stackelberg leader solely interested in minimizing the cost of the flow that she controls, the ability to commit can give an unbounded reduction in the cost of the leader's flow [3].

Various other, similar ratios have been studied, such as the price of stability (which considers the best rather than the worst Nash equilibrium) [1], and analogous concepts using correlated rather than Nash equilibria [4]. However, to our knowledge, beyond the one exception noted above none of these ratios have considered the benefit of commitment.

Our contributions. In the rest of this paper we first formally define the *VoPC*, *VoMC*, and *MvP* ratios, in Section 2. In Section 3, as a "warm-up," we investigate these ratios in 2×2 normal-form and symmetric normal-form games. We investigate these ratios in extensive-form and security games in Section 4. In Section 5, we investigate these ratios in atomic selfish routing games with k-nomial costs (and in Section 5.2 and Section 5.3 we extend these results to symmetric games and arbitrary cost functions). Figure 2 contains a summary of our results. Immediately noticeable is the unboundedness of many of these results. In contrast, in Section 6, we look at experimental results on various distributions over games. Here, we find that generally, the ability to commit to a mixed strategy provides only limited benefits to player 1's utility relative to simultaneous-move equilibrium concepts, though in some cases some of the equilibria are significantly worse for player 1 (particularly, unsurprisingly, ones that are selected to be as bad as possible for player 1). Commitment to a pure strategy is generally almost as beneficial to the leader as commitment to a mixed strategy, with the exception of zero-sum games.

2 Definitions

We now define the concepts of VoPC, VoMC, and MvP formally. The main difficulty in doing so is that we must specify which outcome materializes after player 1 commits, or when player 1 does not commit. That is, what solution concept is used for the remaining game? In Example 1, this was unambiguous: if player 1 commits, player 2 faces a straightforward single-agent optimization problem; if player 1 does not commit, the game is solvable by iterated strict dominance. However, we will not be so lucky for every game: for example, if there are at least three players then even after player 1's commitment the game may not be solvable by iterated strict dominance. (When there are more than two players we assume that only player 1 is able to make a commitment.) Generally, we will have to use a solution concept that is defined everywhere. For example, we can assume that the worst Nash equilibrium for player 1 will result—or the best, or the worst correlated equilibrium, etc. Recall that a Nash equilibrium is a solution concept that assumes that each player knows the equilibrium strategy of the other player, and given this information has no incentive to change their own strategy. An action is said to be a *best response* if under the information a player has available to them, the action they take maximizes their own utility.

Game type	VoP	VoM	MvP
Normal-form games (2×2)	∞ (ISD)	∞ (ISD)	∞ (ISD)
α-close Symmetric normal-form games (3×3)	$\frac{1+3\alpha}{1-\alpha}$(ISD)	$\frac{1+3\alpha}{1-\alpha}$ (ISD)	∞ (ISD)
Symmetric normal-form games (3×3)	∞ (ISD)	∞ (ISD)	∞ (ISD)
Extensive-form games (3 leaves)	∞ (BI)	∞ (BI)	∞ (BI)
Security games (2×2)	∞ (Unique-Correlated)	∞ (Unique-Correlated)	∞ (Unique-Correlated)
Atomic selfish routing			
k-nomial costs (n players)	n^k (ISD)	n^k (ISD)	n^k (ISD)
Symmetric, k-nomial costs (2 players, directed edges)	2^k (Nash)	2^k (Nash)	2^k (Nash)
Arbitrary costs (2 players)	∞ (ISD)	∞ (ISD)	∞ (ISD)

Fig. 2 Summary of results. In parentheses, we give the strongest solution concept for which we prove the lower bound: iterated strict dominance (ISD), backward induction (BI), unique correlated equilibrium, or Nash equilibrium. More details appear in the relevant sections.

Let SC denote a solution concept where both of the players best respond under some information model (although the action they take might not look like a best response under a different information model) and let $u_1(SC(G))$ be the utility that player 1 receives under that solution concept in game G. This includes, but is not limited to: Nash equilibrium (and its various refinements), correlated equilibrium, iiterated strict dominance, and backwards induction. Of course, not all of these solution concepts make sense for any given class ($\mathscr{C}$) of games. Let S_1 (resp. Σ_1) be the set of player 1's pure (resp. mixed) strategies. (Of course, $S_1 \subseteq \Sigma_1$.) Let $\sigma_1 \in \Sigma_1$ be a strategy for player 1; let $G|_{\sigma_1}$ be the game among the remaining players that results after player 1 commits to σ_1. Then the value of pure commitment (w.r.t. SC) is defined as $VoPC^{SC}(\mathscr{C})$:

$$VoPC^{SC}(\mathscr{C}) = \sup_{G\in\mathscr{C}} \frac{\sup_{s_1\in S_1}\{u_1(SC(G|_{s_1}))\}}{u_1(SC(G))}$$

Similarly, we get the following definitions for the value of mixed commitment, $VoMC^{SC}(\mathscr{C})$:

$$VoMC^{SC}(\mathscr{C}) = \sup_{G\in\mathscr{C}} \frac{\sup_{\sigma_1\in \Sigma_1}\{u_1(SC(G|_{\sigma_1})\}}{u_1(SC(G))}$$

and the mixed vs. pure ratio, $MvP^{SC}(\mathscr{C})$:

$$MvP^{SC}(\mathscr{C}) = \sup_{G\in\mathscr{C}} \frac{\sup_{\sigma_1\in \Sigma_1}\{u_1(SC(G|_{\sigma_1})\}}{\sup_{s_1\in S_1}\{u_1(SC(G|_{s_1})\}}$$

Some games are more naturally modeled by cost functions c_i rather than utility functions u_i. In that case, we redefine

$$VoPC^{SC}(\mathscr{C}) = \sup_{G\in\mathscr{C}} \frac{c_1(SC(G))}{\inf_{s_1\in S_1}\{c_1(SC(G|_{s_1}))\}}$$

etc. We are assuming nonnegative utility and cost functions everywhere. In our results below, to simplify the notation, we will omit $\mathscr{C}$ when it is clear from context.

We have the following trivial lemma:

Lemma 1 *For any SC, $VoMC^{SC}(\mathscr{C}) \geq VoPC^{SC}(\mathscr{C})$.*

Proof The sup in the numerator of VoMC is over a larger set than the sup in the numerator of VoPC (in the case of costs, the inf in the denominator of VoMC is over a larger set). ■

Any action that gives less utility than another against all opponent actions is said to be strictly dominated and cannot be a part of any Nash equilibria. Iterated strict dominance (ISD) takes advantage of this fact by removing strictly dominated strategies from the game in an iterative manner, checking to see if new actions can be removed on each iteration until the process reaches convergence. However, for ISD, the expressions above are not well-defined because for any game G that is not solvable by ISD, $u_1(SC(G))$ is not defined. Hence, when we consider $VoPC^{\text{ISD}}$ (etc.), we require that the outer sups are taken only over games solvable by ISD and the inner sups are taken only over strategies that result in a game solvable by ISD. Similarly, we can consider the solution concept UniqueNash, which can be applied only when the game has a unique Nash equilibrium. Because a game solvable by ISD has a unique correlated equilibrium, we get:

Lemma 2 *For any class of games $\mathscr{C}$, for any $SC \in \{\text{BestNash, WorstNash, UniqueNash, BestCorrelated, WorstCorrelated}\}$, we have $VoPC^{\text{ISD}}(\mathscr{C}) \leq VoPC^{SC}(\mathscr{C})$, $VoMC^{\text{ISD}}(\mathscr{C}) \leq VoMC^{SC}(\mathscr{C})$, and $MvP^{\text{ISD}}(\mathscr{C}) \leq MvP^{SC}(\mathscr{C})$.*

As a result, if we can prove a high lower bound for ISD, it immediately implies a high lower bound for these other concepts. Similarly for UniqueNash, we get:

Lemma 3 *For any class of games $\mathscr{C}$, for any $SC \in \{\text{BestNash, WorstNash}\}$, we have $VoPC^{\text{UniqueNash}}(\mathscr{C}) \leq VoPC^{SC}(\mathscr{C})$, $VoMC^{\text{UniqueNash}}(\mathscr{C}) \leq VoMC^{SC}(\mathscr{C})$, and $MvP^{\text{UniqueNash}}(\mathscr{C}) \leq MvP^{SC}(\mathscr{C})$.*

The game of chicken (with the payoffs illustrated in Figure 3) is an example where the value of commitment differs drastically between different solution concepts. The best commitment strategy for player 1 is to commit to a pure strategy of Straight, which gives her a utility of 1. If we compare this to the best Nash equilibrium for player 1, where player 1 plays Straight and player 2 plays Dodge, we find that $VoPC^{\text{BestNash}} = \frac{1}{1} = 1$. However, if we instead compare this to the worst Nash equilibrium for player 1, where player 1 plays Dodge and player 2 plays Straight, we find that $VoPC^{\text{WorstNash}} = \frac{1}{\varepsilon}$, which goes to ∞ as $\varepsilon \to 0$.

	Dodge	Straight
Dodge	(1/2,1/2)	(ε,1)
Straight	(1,ε)	(0,0)

Fig. 3 Chicken

	L	R
U	(ε,1)	(1,0)
D	(0,0)	($1-\varepsilon$,1)

Fig. 4 $VoPC^{\text{ISD}}(NF_{2\times2})$ and $VoMC^{\text{ISD}}(NF_{2\times2})$

	L	R
U	(ε,1)	(1,0)
D	(0,0)	(2ε,1)

Fig. 5 Game for $MvP^{\text{ISD}}(NF_{2\times2})$

2.1 Close-to-Constant-Sum Games

In two-player constant-sum games, by the minimax theorem, each player has a value that she can guarantee herself, and these values sum to the constant. Therefore, the players will obtain these values under any reasonable solution concept, including Stackelberg solutions. Hence, the value of commitment is 1 in such games. This leads to the interesting question of what happens when the game is *close* to constant-sum. We say that a two-player game is α-*close* to a constant sum c if for all outcomes o, $|u_1(o)+u_2(o)-c| \leq \alpha$. We immediately obtain the following lemma:

Lemma 4 *In families of games where one can scale all players' utilities (simultaneously) by any constant and add any constant to a single player's utility function, any game with nonnegative utilities can be made α-close to a constant sum c, for any α and c, without modifying the VoPC, VoMC, and MvP ratios.*

Proof Take an arbitrary game in a family meeting the above qualifications. Let u_i^h be the highest payoff in the game for player i, and let $u^h = \max\{u_1^h, \ldots, u_n^h\}$. Next, we modify each payoff as follows. For each outcome o and for each player $i \in \{1, \ldots, (n)\}$, let $u_i'(o) = \frac{c}{n} + \frac{\alpha}{n \cdot u^h} u_i(o)$. Each modified payoff is in $[\frac{c}{n}, \frac{c}{n} + \frac{\alpha}{n}]$, so this new game is α-close to c. We note that we have applied an affine transformation to each player's utility function, and hence each player's behavior is unchanged. Additionally, for player 1, we have only scaled every payoff by the same amount. Thus, in all three of the ratios of interest, both the numerator and the denominator have been scaled by this value, and hence the ratio has not changed. ■

This result may appear to make the notion of α-closeness to a constant sum c uninteresting in this context, because it suggests that any ratio that can be obtained can also be obtained by a game that is α-close to a constant sum c (for any α and c). This is indeed the case for many of the families that we study, but in one case, namely symmetric normal-form games, the family does not satisfy this condition. We will show that the ratios that can be obtained depends on α for this family of games in Section 3.

3 Normal-Form and Symmetric Normal-Form Games

Let us start with some classes of games for which the ratios are easy to determine.

Normal-form games.

First, consider the class of 2×2 normal-form games ($NF_{2\times 2}$). Here, we find that for $NF_{2\times 2}$:

$$VoPC^{\text{ISD}} = VoMC^{\text{ISD}} = MvP^{\text{ISD}} = \infty.$$

Example 2 For $VoPC^{\text{ISD}} = \infty$, consider the game in Figure 4. In this game, U strictly dominates D, so that the game is solvable by iterated strict dominance, resulting in a utility of ε for player 1. If player 1 commits to the pure strategy D, player 2's best response will be R, resulting in a utility of $1-\varepsilon$ for player 1. Thus $VoPC^{\text{ISD}} \geq \frac{1-\varepsilon}{\varepsilon}$ for any ε, hence by letting $\varepsilon \to 0$, we get $VoPC^{\text{ISD}} = \infty$. $VoMC^{\text{ISD}} = \infty$ follows from $VoPC^{\text{ISD}} = \infty$ and Lemma 1.

For $MvP^{\text{ISD}} = \infty$, consider the game in Figure 5. The optimal pure strategy to commit to is again D, but this time this only gives player 1 a utility of 2ε. However, if player 1 can commit to a mixed strategy, then she can commit to $\frac{1}{2} - \delta\ U$, $\frac{1}{2} + \delta\ D$. Against this mixed strategy, player 2 still prefers to play R, and this way player 1 can get an expected utility arbitrarily close to $1/2+\varepsilon$. Thus $MvP^{\text{ISD}} \geq (1/2+\varepsilon)/(2\varepsilon)$, hence by letting $\varepsilon \to 0$, we get $MvP^{\text{ISD}} = \infty$.

Recall that although we only consider the restrictive solution concept ISD (Iterated Strict Dominance) here, this result also provides a lower bound on $VoPC^{\text{SC}}$, $VoMC^{\text{SC}}$ and MvP^{SC} for $SC \in \{\text{BestNash}, \text{WorstNash}, \text{UniqueNash}, \text{BestCorrelated}, \text{WorstCorrelated}\}$ (Lemma 2).

We can prove a similar result even when the games are α-close to constant sum.

Corollary 1 *For the the class of games $NF\alpha_{2\times 2}$ of 2×2 normal-form games that are α-close to a constant sum of 1:*

$$VoPC^{\text{ISD}} = VoMC^{\text{ISD}} = MvP^{\text{ISD}} = \infty.$$

Proof This follows from Example 2 and Lemma 4. ■

Symmetric Normal-Form Games.

Next, we consider what happens when we further constrain the class of games to be symmetric.

	L	C	R
U	$(\frac{1-\alpha}{2},\frac{1-\alpha}{2})$	$(\frac{1-\alpha}{2}+\varepsilon,\frac{1-\alpha}{2}-\varepsilon)$	$(1+\alpha,0)$
M	$(\frac{1-\alpha}{2}-\varepsilon,\frac{1-\alpha}{2}+\varepsilon)$	$(\frac{1-\alpha}{2},\frac{1-\alpha}{2})$	$(\frac{1+3\alpha}{2}-2\varepsilon,\frac{1-\alpha}{2}+2\varepsilon)$
D	$(0,1+\alpha)$	$(\frac{1-\alpha}{2}+2\varepsilon,\frac{1+3\alpha}{2}-2\varepsilon)$	$(\frac{1}{2},\frac{1}{2})$

Fig. 6 Game for $VoPC^{\text{ISD}}(SNF\alpha_{3\times 3})$ and $VoMC^{\text{ISD}}(SNF\alpha_{3\times 3})$

	L	C	R
U	$(\frac{1}{2},\frac{1}{2})$	(0,1)	(1,0)
M	(1,0)	$(\frac{1}{2},\frac{1}{2})$	(0,1)
D	(0,1)	(1,0)	$(\frac{1}{2},\frac{1}{2})$

Fig. 7 Game for $MvP^{\text{ISD}}(SNF\alpha_{3\times 3})$ (rock-paper-scissors)

Theorem 1 *For the class of two-player symmetric normal-form games that are α-close to a constant sum of 1 ($SNF\alpha$), for any solution concept SC under which both players end up playing a best response:*

$$VoPC^{SC} \leq \frac{1+3\alpha}{1-\alpha} \text{ and } VoMC^{SC} \leq \frac{1+3\alpha}{1-\alpha}.$$

Proof First, we argue that regardless of what the other player does, each player has a response that gives her a utility of at least $\frac{1-\alpha}{2}$. For the sake of contradiction, w.l.o.g., suppose that player 1 has a strategy for which player 2's best response gives him a utility of less than $\frac{1-\alpha}{2}$. By symmetry, this would mean that player 2 also has a strategy for which player 1's best response gives her a utility of less than $\frac{1-\alpha}{2}$. But then, if both players played these strategies, their combined utility would be strictly less than $1-\alpha$, contradicting the game's α-closeness. This shows that the denominator for both of these ratios is at least $\frac{1-\alpha}{2}$. Second, we argue that regardless of how player 1 commits, she can achieve at most $\frac{1+3\alpha}{2}$. Since player 2 is guaranteed to have a response to this commitment that gives him at least $\frac{1-\alpha}{2}$ and because the sum of the two utilities can be at most $1+\alpha$, the most that player 1 can hope to achieve through commitment is $1+\alpha-\frac{1-\alpha}{2}=\frac{1+3\alpha}{2}$. Thus $VoPC^{SC} \leq \frac{1+3\alpha}{1-\alpha}$ and $VoMC^{SC} \leq \frac{1+3\alpha}{1-\alpha}$. ■

If we restrict ourselves to 3×3 symmetric normal-form games that are α-close to a constant sum of 1 for $\alpha < 1$ ($SNF\alpha_{3\times 3}$) we can show a corresponding upper bound:

$$VoPC^{\text{ISD}} = VoMC^{\text{ISD}} \geq \frac{1+3\alpha}{1-\alpha} \text{ and } MvP^{\text{ISD}} = \infty.$$

Example 3 For $VoPC^{\text{ISD}} \geq \frac{1+3\alpha}{1-\alpha}$, consider the game in Figure 6. First, because $\alpha < 1$, M and C are dominated by U and L, respectively. After removing these, B and R are dominated by U and L, respectively. Hence, the game is solvable by ISD, resulting in a utility of $\frac{1-\alpha}{2}$ for player 1. In contrast, if player 1 commits to the pure strategy M, then player 2's best response is R, resulting in a utility of $\frac{1+3\alpha}{2}-2\varepsilon$ for player 1. Thus, $VoPC^{\text{ISD}} \geq \frac{\frac{1+3\alpha}{2}-2\varepsilon}{\frac{1-\alpha}{2}}$ for any $\varepsilon > 0$, hence by letting $\varepsilon \to 0$, we get $VoPC^{\text{ISD}} \geq \frac{1+3\alpha}{1-\alpha}$. $VoMC^{\text{ISD}} \geq \frac{1+3\alpha}{1-\alpha}$ follows from $VoPC^{\text{ISD}} \geq \frac{1+3\alpha}{1-\alpha}$ and Lemma 1.

For $MvP^{\text{ISD}} = \infty$, consider the constant-sum game in Figure 7 (rock-paper-scissors). Any pure strategy that player 1 can commit to has a response by player 2 that gives player 1 a utility of 0. However, if player

1 commits to the mixed strategy (1/3, 1/3, 1/3), then any response by player 2 results in a payoff of 1/2 for player 1. Thus $MvP^{\mathrm{ISD}} = \frac{1/2}{0} = \infty$.

By letting α approach 1, we see that:

$$VoPC^{\mathrm{ISD}}(SNF_{3\times3}) = VoMC^{\mathrm{ISD}}(SNF_{3\times3}) = MvP^{\mathrm{ISD}}(SNF_{3\times3}) = \infty.$$

4 Extensive-Form and Security games

In this section we will look at two additional classes of games.
Extensive-form games.

First we will consider the class $PIEF$ of perfect-information extensive-form games. An extensive-form game gives an explicit ordering of player actions and what information is available to each player at each action. For extensive-form games we will be concerned with comparing with the solution concept of backwards induction (BI). We will now show that for $PIEF$:

$$VoPC^{\mathrm{BI}} = VoMC^{\mathrm{BI}} = MvP^{\mathrm{BI}} = \infty$$

Example 4 For $VoPC^{\mathrm{BI}} = \infty$, consider the game in Figure 8. (It should be emphasized that player 1, the player with commitment power, moves *second* in this game.) In this game, L dominates R for player 1. Given this, at the top level L dominates R for player 2, resulting in a utility of ε for player 1. If player 1 commits to R, then player 2 will prefer to go right, resulting in a utility of $1-\varepsilon$ for player 1. Thus $VoPC^{\mathrm{BI}} \geq \frac{1-\varepsilon}{\varepsilon}$, hence by letting $\varepsilon \to 0$, we get $VoPC^{\mathrm{BI}} = \infty$. As usual, $VoMC^{\mathrm{BI}} = \infty$ follows from $VoPC^{\mathrm{BI}} = \infty$ and Lemma 1.

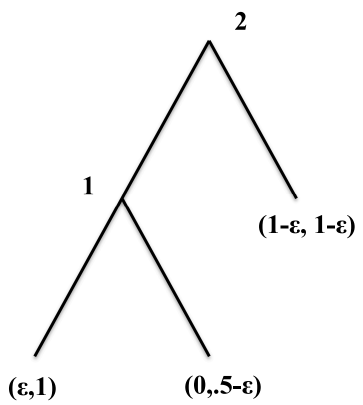

Fig. 8 Game for $VoPC^{\mathrm{BI}}(PIEF)$ and $VoMC^{\mathrm{BI}}(PIEF)$

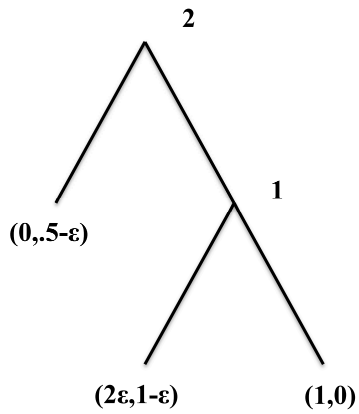

Fig. 9 Game for $MvP^{\mathrm{BI}}(PIEF)$

For $MvP^{\mathrm{BI}} = \infty$, consider the game in Figure 9. The optimal strategy to commit to is L, leading to a utility of 2ε for player 1. If player 1 commits to a mixed strategy of $\frac{1}{2}$ R, $\frac{1}{2}$ L, then player 2 will still prefer to go right. In this case player 1's expected utility is $\frac{1-\varepsilon}{2}$. Thus $MvP^{\mathrm{BI}} \geq ((1-\varepsilon)/2)/\varepsilon$, hence by letting $\varepsilon \to 0$, we get $MvP^{\mathrm{BI}} = \infty$.

We also have the following corollary for when the game is restricted to be α-close.

Corollary 2 *For perfect-information extensive-form games that are α-close to constant sum ($PIEF\alpha$):*

$$VoPC^{\mathrm{BI}} = VoMC^{\mathrm{BI}} = MvP^{\mathrm{BI}} = \infty.$$

Proof This follows from Example 4 and Lemma 4. ■

Security games.

Motivated by the various applications of computing optimal mixed strategies to commit to in security domains (as discussed in the introduction), we now consider a small subclass of the class of security games defined by Kiekintveld et al. [10]. We will call this class *simple security games* (*SSG*). In a simple security game, there is a set of *targets* T. The defender (player 1) chooses a target to defend; the attacker (player 2) chooses a target to attack. Each player's utility depends on two things: (1) which target the attacker attacks, and (2) whether the defender has chosen to defend that target. Thus, $u_1^c(t)$ (resp., $u_2^c(t)$) is the defender's (resp., attacker's) utility when target t is attacked and it is defended, and $u_1^u(t)$ (resp., $u_2^u(t)$) is the defender's (resp., attacker's) utility when target t is attacked and it is not defended. We require that for any t, $u_1^c(t) > u_1^u(t)$ and $u_2^c(t) < u_2^u(t)$.

	T_1	T_2
T_1	$(\varepsilon,0)$	$(\frac{1}{2},1)$
T_2	$(0,1)$	$(1,0)$

Fig. 10 Game for $VoPC^{\text{UniqueCorrelated}}(SSG_{2\times 2})$ and $VoMC^{\text{UniqueCorrelated}}(SSG_{2\times 2})$

	T_1	T_2
T_1	(1,0)	(0,1)
T_2	(0,1)	(1,0)

Fig. 11 Game for $MvP^{\text{UniqueCorrelated}}(SSG_{2\times 2})$ (matching pennies)

We will now show that for 2×2 simple security games ($SSG_{2\times 2}$):

$$VoPC^{\text{UniqueCorrelated}} = VoMC^{\text{UniqueCorrelated}} = MvP^{\text{UniqueCorrelated}} = \infty.$$

Example 5 For $VoPC^{\text{UniqueCorrelated}} = \infty$, consider the game in Figure 10. Consider the following profile of mixed strategies: player 1 plays $\frac{1}{2}$ T_1 and $\frac{1}{2}$ T_2; player 2 plays ($\frac{1}{1+2\varepsilon}$) T_1, $\frac{2\varepsilon}{1+2\varepsilon}$ T_2 for player 2). It is not hard to see that this is the unique Nash equilibrium of the game; we will now prove the stronger claim that it is even the unique correlated equilibrium of the game. Let p_{ij} (for $i, j \in \{1,2\}$) be the probability that player 1 plays T_i and player 2 plays T_j. Then, the correlated equilibrium constraints are

- $p_{21} \geq p_{11}$ (when player 2 is recommended to play T_1)
- $p_{12} \geq p_{22}$ (when player 2 is recommended to play T_2)
- $\varepsilon p_{11} + (1/2)p_{12} \geq p_{12} \Leftrightarrow 2\varepsilon p_{11} \geq p_{12}$ (when player 1 is recommended to play T_1)
- $p_{22} \geq \varepsilon p_{21} + (1/2)p_{22} \Leftrightarrow p_{22} \geq 2\varepsilon p_{21}$ (when player 1 is recommended to play T_2)

From this, we obtain $p_{22} \geq 2\varepsilon p_{21} \geq 2\varepsilon p_{11} \geq p_{12} \geq p_{22}$, so it follows that these four quantities must all be the same. Combined with the constraint that the probabilities sum to 1, we obtain the unique solution $p_{11} = p_{21} = \frac{1}{2(1+2\varepsilon)}$ and $p_{12} = p_{22} = \frac{2\varepsilon}{2(1+2\varepsilon)}$, which is consistent with the mixed strategies above. This equilibrium leads to an expected utility for player 1 of $\frac{2\varepsilon}{1+2\varepsilon} < 2\varepsilon$.

On the other hand, if player 1 commits to the pure strategy T_1, player 2's best response is T_2, so that player 1's utility is $\frac{1}{2}$. Thus, $VoPC^{\text{UniqueCorrelated}} \geq (1/2)/(2\varepsilon)$, hence by letting $\varepsilon \to 0$, we get $VoPC^{\text{UniqueCorrelated}} = \infty$. $VoMC^{\text{UniqueCorrelated}} = \infty$ follows from $VoPC^{\text{UniqueCorrelated}} = \infty$ and Lemma 1.

For $MvP^{\text{UniqueCorrelated}} = \infty$, consider the game in Figure 11 (matching pennies). Any pure strategy that player 1 can commit to has a response by player 2 that gives player 1 a utility of 0. However, if player 1 commits to the mixed strategy (1/2, 1/2), then any response by player 2 results in a payoff of 1/2 for player 1. Thus $MvP^{\text{UniqueCorrelated}} = \infty$.

Korzhyk et al. [13] show that in security games (satisfying a minor assumption[2]), the optimal mixed strategy to commit to for the defender is always a Nash equilibrium strategy as well. This may seem to contradict our result that commitment can be of value in these games, but in fact there is no contradiction. The optimal thing to do when committing to a mixed strategy is to commit to a strategy that is very close to the equilibrium strategy, but incentivizes the attacker to play a response that is much better for the defender than the attacker equilibrium strategy. Also, in that paper the attacker breaks ties in the defender's favor, so that the optimal mixed strategy to commit to is actually exactly the defender's equilibrium strategy.

Again, we have a matching corollary for α-close simple security games.

Corollary 3 *For* 2×2 *α-close simple security games (SSG$\alpha_{2\times 2}$):*

$$VoPC^{\text{UniqueCorrelated}} = VoMC^{\text{UniqueCorrelated}} = MvP^{\text{UniqueCorrelated}} = \infty.$$

Proof This follows from Example 5 and Lemma 4. ■

5 Routing games

In this section, we present our technically most interesting results. We will consider what are known as *atomic selfish routing* games [8]. In an atomic selfish routing game, each player has a pair of nodes, a *source* node and a *target* node, and is interested in routing a given amount of indivisible flow from the former to the latter. The cost incurred by each player depends both on the path chosen and the paths that the other players choose. Since a follower will never reuse an edge in any undominated solution, we can prove finite upper bounds on the amount of utility that can be gained by using commitment to manipulate the follower(s) in many of our routing settings.

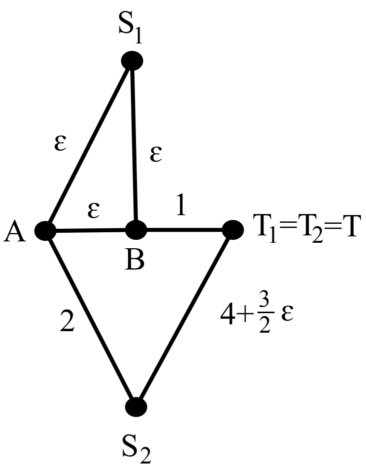

Fig. 12 2-player *LAR* routing game for *VoPC* and *VoMC*

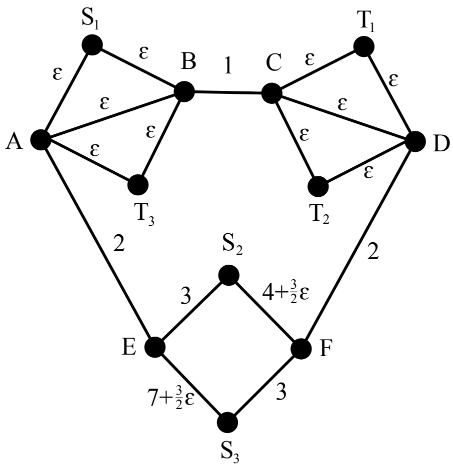

Fig. 13 3-player *LAR* routing game for *VoPC* and *VoMC*

[2]They assume that every subset of a valid schedule is also a valid schedule.

5.1 k-nomial Atomic Routing on Arbitrary Graphs

We first focus on the case where each player controls an indivisible unit amount of traffic, the edges are undirected, and for every edge, the cost of using that edge scales as a general monomial, x^k, with the number of players using that edge. That is, each edge e has a coefficient c_e, and if n_e players use that edge, each of them incurs a cost of $c_e n_e^k$. We refer to this class of games as *kAR* (k-nomial Atomic Routing). We use the cost-based versions of the definitions of *VoPC*, *VoMC*, and *MvP*.

Theorem 2 *For k-nomial Atomic Routing with n players ($kAR_{n\text{ player}}$):*

$$VoPC^{\mathrm{SC}} \leq n^k,\ VoMC^{\mathrm{SC}} \leq n^k,\ \text{and}\ MvP^{SC} \leq n^k$$

for any solution concept SC where player 1 maximises her utility under the information available to her and the other players do not play dominated strategies.

Proof For any given routing game G, we can remove all the players but player 1 to obtain a game G^1. In this game, player 1 will just choose the lowest-cost path. That same path is still an option in G; in the worst case, each of the other $n-1$ players uses all the edges on that path (but each of them will use each edge at most once, because using an edge more than once is a dominated strategy). Because player 1 best-responds under SC, she must do at least that well; hence, we have $c_1(SC(G)) \leq n^k c_1(SC(G^1))$. Moreover, no matter which strategy σ_1 player 1 commits to, she can do no better than $c_1(SC(G^1))$; hence, $c_1(SC(G|_{\sigma_1})) \geq c_1(SC(G^1))$. It follows that for any game $n^k c_1(SC(G|_{\sigma_1})) \geq c_1(SC(G))$, hence $VoMC^{SC} \leq n^k$ and $VoPC^{SC} \leq n^k$. (The preceding discussion assumes that SC can be used to solve G and $G|_{\sigma_1}$, which may not be the case if, for example, $SC = \mathrm{ISD}$, but this does not affect upper bounds on the ratio.) In fact, from this we can also conclude that $MvP^{SC} \leq n^k$: under commitment to a pure strategy, at the very least player 1 can commit to the path s_1 she would choose in G^1, so that $c_1(SC(G|_{s_1})) \leq n^k c_1(SC(G^1))$. Then, for any σ_1, $n^k c_1(SC(G|_{\sigma_1})) \geq n^k c_1(SC(G^1)) \geq c_1(SC(G|_{s_1}))$, so $MvP^{SC} \leq n^k$. ■

We next consider a pair of games that give some insight in how the construction of the proof of the upper bound will proceed. To keep things simple, for now we will assume that that the cost of using an edge scales linearly with the number of players using routing flow over it ($k = 1$). We will refer to this as a Linear Atomic Routing (LAR) game. We start with a two-player game that gives the desired VoPC value of 2.

Example 6 For Linear Atomic Routing with two players ($LAR_{2\text{ player}}$), the game depicted in Figure 12 has $VoPC^{\mathrm{ISD}} \geq 2$.

Without commitment, this routing game is solvable by ISD, as follows. First of all, it is easy to see that any strategy that visits the same vertex twice is dominated, so we can remove such strategies. After that, it is a dominant strategy for player 1 to take the path $S_1 \to B \to T$. Given this path for player 1, the best-response strategy for player 2 is to take the path $S_2 \to A \to B \to T$ (at a cost of $4+\varepsilon$). This ISD solution gives player 1 a cost of $2+\varepsilon$.

When we consider pure-strategy commitment, if player 1 commits to the path $S_1 \to A \to B \to T$, then player 2's best response to is to take the alternate path directly from $S_2 \to T$ (because other paths cost at least $4+2\varepsilon$). This results in a cost of only $1+2\varepsilon$ for player 1. It follows that $VoPC^{\mathrm{ISD}} \geq \frac{2+\varepsilon}{1+2\varepsilon}$, hence by letting $\varepsilon \to 0$, we get $VoPC^{\mathrm{ISD}} \geq 2$.

As with the examples in the previous sections, the leader is able to commit to a dominated strategy to influence the choice of route by the follower. Not surprisingly, it turns out that she can gain the most utility when paying a cost of ε is sufficient to discourage the follower from using any of the edges she wishes to use. Next, let us consider how we might add a third player while still achieving the desired VoPC value of 3.

Example 7 For Linear Atomic Routing with three players ($LAR_{3\text{ player}}$), the game depicted in Figure 13 has $VoPC^{\text{ISD}} \geq 3$.

Again, after removing the strategies that visit a node more than once, player 1 has a dominant strategy, namely $S_1 \to B \to C \to T_1$. Once we fix this strategy for player 1, player 3 has a dominant strategy, namely $S_3 \to F \to D \to C \to B \to T_3$ (i.e., this is optimal regardless of what player 2 does; it will always cost player 3 at most $9+2\varepsilon$ to use this path. In contrast, it will cost player 3's at least $9+\frac{5}{2}\varepsilon$ on any path through E). Once we fix this strategy, player 2 has a best response of $S_2 \to E \to A \to B \to C \to T_2$. This ISD solution gives player 1 a cost of $3+2\varepsilon$.

When we consider pure strategy commitment, if player 1 commits to the path $S_1 \to A \to B \to C \to D \to T_1$, then player 2 has a dominant strategy $S_2 \to F \to D \to T_2$ (i.e., this is optimal regardless of what player 3 does; it will always cost player 2 at most $8+\frac{5}{2}\varepsilon$ to use this path. In contrast, it will cost player 2 at least $8+3\varepsilon$ on any path through E). Once we fix this strategy, player 3 has a best response of $S_3 \to E \to A \to T_3$. This results in a cost of $1+4\varepsilon$ for player 1. It follows that $VoPC^{\text{ISD}} \geq \frac{3+2\varepsilon}{1+4\varepsilon}$, hence by letting $\varepsilon \to 0$, we get $VoPC^{\text{ISD}} \geq 3$.

Note that in this construction player 1 is again paying a minimal cost, in this case to influence player 2, and in addition player 2's response has a cascading effect of also influencing player 3 to change routes as well. We will now leverage this technique to prove the following lower bound:

Theorem 3 *For k-nomial Atomic Routing with n players:*

$$VoPC^{\text{ISD}}(kAR_{\text{n player}}) \geq n^k,\ VoMC^{\text{ISD}}(kAR_{\text{n player}}) \geq n^k,\ \text{and}\ MvP^{\text{ISD}}(kAR_{\text{n player}}) \geq n^k.$$

Proof We will start by proving that $VoPC^{\text{ISD}}(kAR_{\text{n player}}) \geq n^k$ (which immediately implies $VoMC^{\text{ISD}}(kAR_{\text{n player}}) \geq n^k$ by Lemma 1).

First, let us start by relaxing the requirement that $k=1$ from the three-player case discussed in Example 7. This will result in the following changes to the construction. First, we need to increase the costs of the edges between the upper and lower sections, to at least 3^k-1. Second, we also need to make some corresponding adjustments to the coefficients in the lower section to obtain similar strategic effects. Figure 14 gives the updated construction for the three-player (two-follower) case for general k.

We now use this game to show that $VoPC^{\text{ISD}}(kAR_{3\text{ player}}) \geq 3^k$. Again, after removing the obviously dominated strategies that visit a node more than once, player 1 has a dominant strategy of $S_1 \to B \to C \to T_1$. Once we fix this strategy for player 1, player 3 has a dominant strategy, namely $S_3 \to F \to D \to C \to B \to T_3$. Once we fix this, player 2 has a best response of $S_2 \to E \to A \to B \to C \to T_2$. This ISD solution gives player 1 a cost of $3^k+2\varepsilon$.

When we consider pure strategy commitment, if player 1 commits to the path $S_1 \to A \to B \to C \to D \to T_1$, then player 2 has a dominant strategy of $S_2 \to F \to D \to T_2$. This leads to player 3 having a best response of $S_3 \to E \to A \to T_3$, and results in a cost of $1+4\varepsilon$ for player 1. It follows that $VoPC^{\text{ISD}}(kAR_{3\text{player}}) \geq \frac{3^k+2\varepsilon}{1+4\varepsilon}$ for any ε, hence by letting $\varepsilon \to 0$, we get $VoPC^{\text{ISD}}(kAR_{3\text{ player}}) \geq 3^k$.

The game for the four-player case for $VoPC^{\text{ISD}}(kAR_{4\text{ player}})$ is pictured in Figure 15. Without commitment, the ISD solution is as follows: it is dominant for player 1 to choose $S_1 \to B \to E \to T_1$, then for player 4 to choose $S_4 \to G \to C \to B \to E \to T_4$, then for player 3 to choose $S_3 \to J \to D \to E \to B \to T_3$, and finally for player 2 to choose $S_2 \to H \to A \to B \to E \to T_2$. As a result, player 1's cost is $4^k+2\varepsilon$. On the other hand, if player 1 can commit to a pure strategy, then she can commit to the path $S_1 \to A \to B \to C \to S_1 \to B \to E \to T_1$. it then becomes dominant for player 2 to choose $S_2 \to J \to D \to T_2$, then for player 3 to choose $S_3 \to G \to C \to T_3$, and finally for player 4 to choose $S_4 \to I \to F \to T_4$. As a result, player 1's cost is $1+5\varepsilon$. We can conclude that $VoPC^{\text{ISD}}(kAR_{4\text{ player}}) \geq 4^k$.

The case of general n continues this pattern (note that a n-player game has $n-1$ followers). The construction can be broken into two sections, an upper section that contains the path for player 1 and the targets of the followers, and a bottom section that contains the sources of the followers.

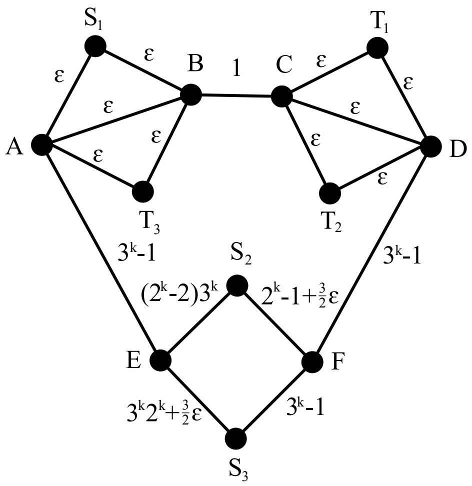

Fig. 14 3-player *kAR* game for *VoPC* and *VoMC*

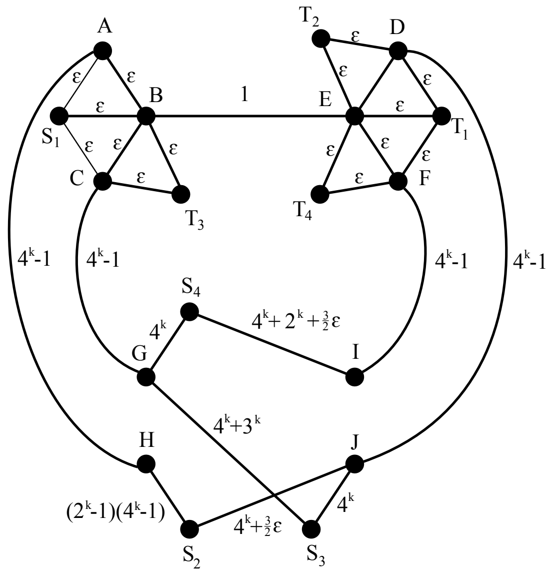

Fig. 15 4-player *kAR* game for *VoPC* and and *VoMC*

Let us first describe the structure of the upper section. We start with the source S_1 and target T_1 nodes for player 1. These two nodes are connected by a path of length 3, with edges of weight ε, 1 and ε (for example, the path $S_1 \to B \to E \to T_1$ in Figure 15). Let us call the two additional nodes on this path u_S and u_T, with u_S being the node adjacent to S_1. Next, for each odd-numbered follower, we add a pair of nodes, one connected to S_1 and u_S by edges of cost ε, the other connected to T_1 and u_T (for example, nodes A and D in Figure 15). Let us refer to these nodes by the number of the (odd-numbered) follower that we add them for—for example, the two nodes added for follower 1 would be u_1 and u_2, with u_1 adjacent to S_1; the ones added for follower 3 would be u_3 and u_4, etc. Finally, for each follower i besides the last, we add a target node T_{i+1}, connected by edges of cost ε to u_{i+1} and u_S if i is even, otherwise to u_{i+1} and u_T. For the last follower $n-1$, we add a target node T_n. If there are an even number of followers, then we connect this node to u_1 and u_S, otherwise to u_n and u_T.

Next, let us describe the structure of the lower section. For the first follower, we add three nodes. The first of these nodes is the source for the first follower, S_2. We add two additional nodes connected to S_2, which we call l_1 and l_2. The edge between l_1 and S_2 has a cost of $(2^k-1)(n^k-1)$, and the edge between l_2 and S_2 has a cost of $n^k + \frac{3}{2}\varepsilon$. For each follower i from 2 to $n-2$, we add two nodes, namely the follower's source S_{i+1}, and l_{i+1}. We also add two edges, one between l_i and S_{i+1} with a cost of n^k, and the other between l_{i+1} and S_{i+1}, with a cost of $n^k + (n-i+1)^k$. If there are an odd number of followers, then for the $(n-1)$th follower we follow a similar pattern, adding two nodes S_n and l_n. Again, we also add two edges, one between l_{n-1} and S_n with a cost of n^k, and one between l_n and S_n with an edge cost of $4^k + 2^k + \frac{3}{2}\varepsilon$. If instead we have an even number of followers, then we only add one node, S_n. However, we still add two edges, one between l_{n-1} and S_n with a cost of n^k, and one between l_n and S_n, with an edge cost of $4^k + 2^k + \frac{3}{2}\varepsilon$. Finally, we connect the upper and lower sections with edges between each pair (l_j, u_j), with an edge cost of $n^k - 1$.

The upper section is designed so that player 1 has a dominant strategy, namely taking the path $S_1 \to u_S \to u_T \to T_1$. If we fix this path for player 1, then player n will have a dominant strategy that involves going through the edge between u_S and u_T; fixing that, the same becomes true for player $n-1$; etc. This results in a cost of $n^k + 2\varepsilon$ for player 1. However, the upper section also allows player 1 to commit to a pure strategy that involves adding u_1 and u_n to her path (visiting S_1 and u_S twice if there an odd number of followers). If we fix this path for player 1, then player 2 will have a dominant strategy that avoids going through the edge between u_S and u_T; fixing that, the same becomes true for player 3, etc. This results in a cost of $1+4\varepsilon$ or $1+5\varepsilon$ (depending on whether there are an even or an odd number of followers) for player 1. It follows that $VoPC^{\text{ISD}}(kAR_{n\text{ player}}) \geq \frac{n^k+2\varepsilon}{1+5\varepsilon}$, hence by letting $\varepsilon \to 0$, we get $VoPC^{\text{ISD}}(kAR_{n\text{ player}}) \geq n^k$. Again, $VoMC^{\text{ISD}}(kAR_{n\text{ player}}) \geq n^k$ follows from $VoPC^{\text{ISD}}(kAR_{n\text{ player}}) \geq n^k$ and Lemma 1.

Finally, we need to show $MvP^{\text{ISD}}(kAR_{n\text{ player}}) \geq n^k$. The construction for $MvP^{\text{ISD}}(kAR_{n\text{ player}})$ is similar to the one for $VoPC^{\text{ISD}}(kAR_{n\text{ player}})$; to illustrate what needs to be modified, we highlight the differences with the three-player (two-follower) case. Let us consider the game in Figure 16. This game is identical to the one in Figure 14, except for some changes to the edge costs. The major change is in the upper section of the graph, where all of the coefficients of ε have been replaced with coefficients of $\frac{3^k-1}{2}$, with the exception of the ones on the most direct path for player 1. Second, there are minor adjustments to the coefficients in the bottom section, to compensate for the changes in the top section and ensure similar comparisons as in the original version.

The effect of the changes is that commitment to a pure strategy is unable to help player 1 relative to taking the direct path (which results in a cost of $3^k+2\varepsilon$). For example, if player 1 commits to the path $S_1 \to A \to B \to C \to D \to T_1$, then players 2 and 3 will not take the edge $B \to C$, but this leaves player 1 with a cost of $2 \cdot 3^k - 1$. She can also commit to the path $S_1 \to A \to B \to C \to T_1$, in which case player 2 will not take the edge $B \to C$, but player 3 will, so that player 1 ends up with a cost of $3^k + 2^k + \varepsilon - 1$. In contrast, if player 1 can commit to a mixed strategy, she can randomize between using the path extended on the left side ($S_1 \to A \to B \to C \to T_1$) 2ε of the time, the path extended on the right side ($S_1 \to B \to C \to D \to T_1$) 2ε of the time, and the shortest path ($S_1 \to B \to C \to T_1$) the rest of the time. This is enough to incentivize both of the followers to not use the edge between B and C, causing the expected cost for player 1 to be $4\varepsilon(3^k+\varepsilon)+(1-4\varepsilon)(1+2\varepsilon) = 1+(4\cdot 3^k-2)\varepsilon - 4\varepsilon^2$. Letting $\varepsilon \to 0$, we obtain $MvP^{\text{ISD}}(kAR_{3\text{ player}}) \geq 3^k$. This modification can be generalized to the $VoPC^{\text{ISD}}(kAR_{n\text{ player}})$ construction for any n, giving us $MvP^{\text{ISD}}(kAR_{n\text{ player}}) \geq n^k$. ■

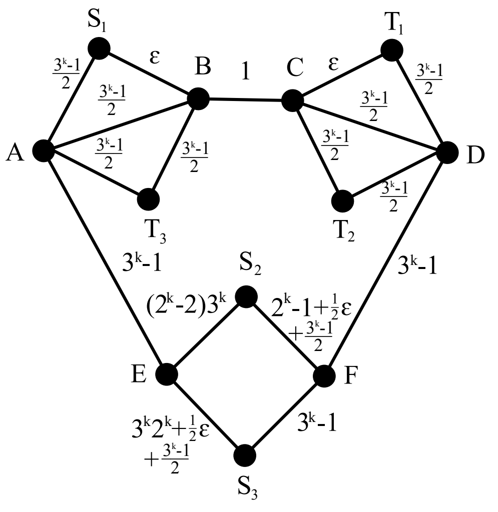

Fig. 16 3-player *kAR* game for *MvP*

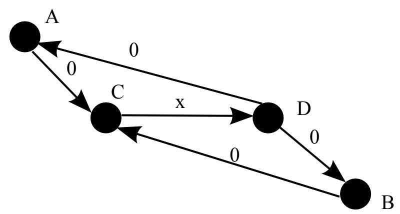

Fig. 17 Gadget for directed edges

Let us next consider the case where we have directed rather than undirected edges. The upper bound in Theorem 2 still holds (the proof did not rely on edges being undirected). On the other hand, the proof of the lower bound in Theorem 3 did rely on the players of the game being able to cross an edge in both directions, so it is not immediately obvious that this lower bound still holds. It turns out, however, that we can simulate an undirected edge using a small number of directed edges.

Lemma 5 *For k-nomial Atomic Routing, it is possible with only a linear (in the number of edges) increase in the number of nodes and edges to transform any undirected graph into a corresponding directed graph with the same VoPC, VoMC and MvP values.*

Proof To perform this transformation we will use the gadget given in Figure 17 to replace each undirected edge. Nodes A and B correspond to the two endpoints of the original undirected edge, and the directed edge from C to D is assigned the cost of the original undirected edge. It is straightforward to check that all traffic across the gadget from A to B as well as from B to A must cross the edge from C to D, thus each of these players incurs the same cost as in the undirected case. As we add 2 nodes and 5 directed edges for each undirected edge in the initial game, this increases the number of nodes and edges by a linear factor. ∎

We can immediately conclude that the same bounds from above hold when we have directed edges. *kARD* (*k*-nomial Atomic Routing with Directed edges).

Theorem 4 *For k-nomial Atomic Routing with Directed edges and n players ($kARD_{n\text{ player}}$):*

$$VoPC^{\mathrm{SC}} \leq n^k,\ VoMC^{\mathrm{SC}} \leq n^k,\ \text{and } MvP^{SC} \leq n^k$$

for any solution concept SC where player 1 *maximises her utility under the information available to her and the other players do not play dominated strategies.*

Theorem 5 *For k-nomial Atomic Routing with Directed edges and n players ($kARD_{n\text{ player}}$):*

$$VoPC^{\mathrm{ISD}} \geq n^k,\ VoMC^{\mathrm{ISD}} \geq n^k,\ \text{and } MvP^{\mathrm{ISD}} \geq n^k.$$

5.2 Symmetric Routing Games

Let us consider what happens if we restrict ourselves to symmetric games, where all players share the same source and target nodes. We consider both the directed and undirected cases here. We will first look at problem that we will call *SkARD* (Symmetric *k*-nomial Atomic Routing with Directed edges) in the two-player (one follower) setting.

Theorem 6 *For Symmetric k-nomial Atomic Routing with Directed edges and* 2 *players ($SkARD_{2\text{ player}}$):*

$$VoPC^{\mathrm{SC}} = VoMC^{\mathrm{SC}} = MvP^{SC} \leq 2^k$$

for any solution concept SC where player 1 *maximises her utility under the information available to her and the other player does not play dominated strategies.*

Proof Because *SkARD* is a special case of *kARD*, upper bounds from the latter still hold, giving us $VoMC^{SC} \leq 2^k$, $VoPC^{SC} \leq 2^k$, and $MvP^{SC} \leq 2^k$.

We now turn to the more difficult part: proving that $VoPC^{SC} \geq 2^k$ (which immediately implies $VoMC^{SC} \geq 2^k$ by Lemma 1). For this bound, let us consider the game in Figure 18.

First, let us consider this game without commitment. It is easy to see that any strategy that visits the same vertex twice is dominated and hence will not be played under *SC*, so we can remove such strategies from consideration. Then, at least one of the paths $S \rightarrow A \rightarrow C$ and $S \rightarrow B \rightarrow C$ will be used with probability at most $1/2$ by player 1. Thus, player 2 always has a strategy that includes the edge

$C \to T$ and that shares the path between S and C with player 1 with probability at most $1/2$, for a total cost of at most $2^k + (1/2)2\varepsilon + (1/2)2^{k+1}\varepsilon = 2^k + (2^k+1)\varepsilon$. Player 2 will prefer this strategy to the alternate direct edge between S and T, and as a result player 1 will share the $C \to T$ edge with player 2, resulting in a cost of at least 2^k for player 1.

When we consider pure-strategy commitment, if player 1 commits to the path $S \to A \to C \to S \to B \to C \to T$, then the path $S \to T$ becomes the unique best response for player 2 (because either path involving the $C \to T$ edge will cost him $2^k + 2^{k+1}\varepsilon$). This is to the benefit of player 1, who no longer has to share the $C \to T$ edge with player 2 and hence has a cost of only $1+5\varepsilon$. It follows that $VoPC^{SC} \geq \frac{2^k}{1+5\varepsilon}$, and hence by letting $\varepsilon \to 0$, we get $VoPC^{SC} = 2^k$. Again, $VoMC^{SC} = 2^k$ follows from $VoPC^{SC} \geq 2^k$ and Lemma 1.

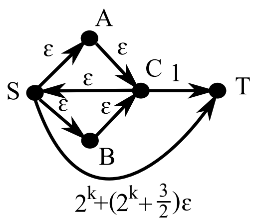

Fig. 18 2-player *SkARD* game for *VoPC* and *VoMC*

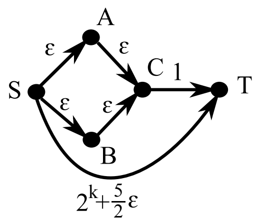

Fig. 19 2-player *SkARD* game for *MvP*

Finally, we need to show that $MvP^{SC} \geq 2^k$. For this proof we will use the game in Figure 19.

Let us first consider the best pure strategy commitment for player 1. The best pure-strategy commitments player 1 has are the path $S \to A \to C \to T$ and the path $S \to B \to C \to T$. (Unlike in the previous example, there is no edge back from $C \to S$, so player 1 cannot commit to use both of the paths between S and C.) In response, player 2 will take the other one of these two paths (which results in a cost of $2^k+2\varepsilon$), and thus player 1's cost will also be $2^k+2\varepsilon$. However, if player 1 is able to commit to a mixed strategy, she can commit to the following: 50% $S \to A \to C \to T$ and 50% $S \to B \to C \to T$. Then, player 2's expected cost for either of the top two paths is $2^k + (1/2)2\varepsilon + (1/2)2^{k+1}\varepsilon = 2^k + (2^k+1)\varepsilon$, and player 2 will prefer to take the direct path $S \to T$. This results in a cost for player 1 of $1+2\varepsilon$. Letting ε approach 0, we obtain $MvP^{\mathrm{SC}} = 2^k$. ■

Let us now consider symmetric atomic routing games with undirected edges, which we will refer to as *SkAR* (Symmetric *k*-nomial Atomic Routing). Below is our result in the two-player (one follower) setting. Our results for *VoPC* and *VoMC* closely resemble the corresponding results for *SkARD*, but at this time we can only show a weaker lower bound of $\frac{1+2^k}{2}$ for *MvP*.

Theorem 7 *For Symmetric k-nomial Atomic Routing with* 2 *players* ($SkAR_{2\ \mathrm{player}}$)*:*

$$VoPC^{\mathrm{SC}} = VoMC^{\mathrm{SC}} = 2^k \text{ and } \frac{1+2^k}{2} \leq MvP^{SC} \leq 2^k$$

for any solution concept SC where player 1 *maximises her utility under the information available to her and the other player does not play dominated strategies.*

Proof As with the previous proof, because *SkAR* is a special case of *kAR*, upper bounds from the latter still hold, giving us $VoMC^{SC} \leq 2^k$, $VoPC^{SC} \leq 2^k$, and $MvP^{SC} \leq 2^k$.

Proving that $VoPC^{SC} \geq 2^k$ (which immediately implies $VoMC^{SC} \geq 2^k$ by Lemma 1) is again the more difficult part. Consider the game in Figure 20.

With similar logic as the previous proof, we conclude that there are three possible paths for player 1 ($S \to A \to C$, $S \to B \to C$ and $A \to C$) and one of the paths will be used with probability at most $1/3$ by player 1. Thus, we can only guarentee that player 2 shares the path between S and C with player 1 with probability at most $1/3$, for a total cost of at most $2^k + (2/3)2\varepsilon + (1/3)2^{k+1}\varepsilon = 2^k + \frac{2(2^k+2)}{3}\varepsilon$. Player 2 will again prefer this strategy to the alternate direct edge between S and T, and as a result player 1 will share the $C \to T$ edge with player 2, resulting in a cost of at least 2^k for player 1.

If player 1 commits to the path $S \to A \to C \to S \to B \to C \to T$, then the path $S \to T$ becomes the unique best response for player 2 (because any path involving the $C \to T$ edge will cost him $2^k + 2^{k+1}\varepsilon$). This causes player 1's cost with commitment to be $1+6\varepsilon$. It follows that $VoPC^{SC} \geq \frac{2^k}{1+6\varepsilon}$, hence by letting $\varepsilon \to 0$, we get $VoPC^{SC} = 2^k$. Again, $VoMC^{SC} = 2^k$ follows from $VoPC^{SC} = 2^k$ and Lemma 1.

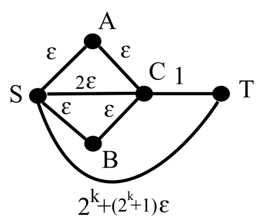

Fig. 20 2-player *SkAR* game for *VoPC* and *VoMC*

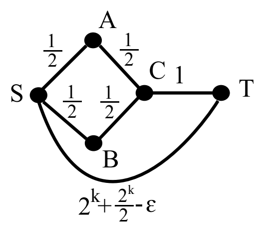

Fig. 21 2-player *SkAR* game for *MvP*

Finally, we need to show $MvP^{\text{BestNash}} \geq \frac{1+2^k}{2}$. Consider the game in Figure 21. As we have undirected edges, we are forced to use significantly larger edge costs on the edges between S, A, B, and C compared to Figure 19.

The best pure-strategy commitments player 1 has are the path $S \to A \to C \to T$ and the path $S \to B \to C \to T$. (Unlike in the previous example, there is an even number of paths between S and C, so player 1 cannot commit to use both of the paths between S and C without using one of the paths twice.) In response, player 2 will take the other out of these two paths, resulting in a cost of $2^k + 1$ for both of the players. However, if player 1 is able to commit to a mixed strategy, she can commit to the following: 50% $S \to A \to C \to T$ and 50% $S \to B \to C \to T$. Then, player 2's expected cost for either of the top two paths is $2^k + (1/2) + (1/2)2^k = 2^k + \frac{2^k+1}{2}$, and player 2 will prefer to take the direct path $S \to T$. This results in a cost for player 1 of 2. Letting ε approach 0, we obtain $MvP^{\text{SC}} \geq \frac{1+2^k}{2}$. ■

5.3 Routing Games with Arbitrary Cost Functions

Let us consider the case where each edge has an arbitrary monotonically increasing cost based on the number of players using the edge. We will refer to this class of games as *AAR* (Arbitrary Atomic Routing). Here, we get arbitrarily large ratios even with two players.

Theorem 8 *For Arbitrary Atomic Routing with 2 players ($AAR_{2\text{ player}}$):*

$$VoPC^{\text{ISD}} = VoMC^{\text{ISD}} = MvP^{\text{ISD}} = \infty.$$

Proof First, let us prove that $VoPC^{\text{ISD}} = \infty$. Consider the game in Figure 22, where all edges but the edge between $B \to T$ have linear costs (indicated with a coefficient as before), and $B \to T$ has a cost of ε if one person uses it, but a cost of 1 if both players use it.

Again, after removing the obvously dominated strategies that visit a node more than once, player 1 has a dominant strategy of $S_1 \to B \to T$. Given this path for player 1, the best-response strategy for player 2 is to take the path $S_2 \to A \to B \to T$. This ISD solution gives player 1 a cost of $1+\varepsilon$.

When we consider pure strategy commitment, if player 1 commits to the path $S_1 \to A \to B \to T$, then player 2's best response to is to take the alternate direct path $S_2 \to T$. This results in a cost of $1+2\varepsilon$ for player 1. It follows that $VoPC^{\text{ISD}} \geq \frac{1+\varepsilon}{2\varepsilon}$, hence by letting $\varepsilon \to 0$, we get $VoPC^{\text{ISD}} \geq \infty$. $VoMC^{\text{ISD}} \geq \infty$ follows from $VoPC^{\text{ISD}} \geq \infty$ and Lemma 1.

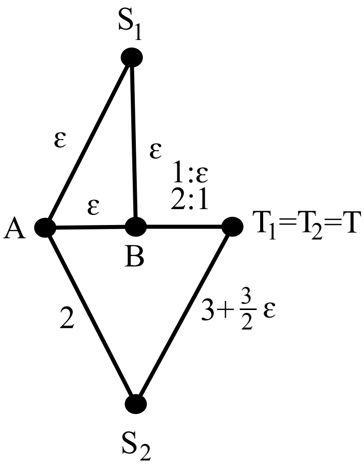

Fig. 22 2-player $AAR_{2\,\text{player}}$ game for *VoPC* and *VoMC*

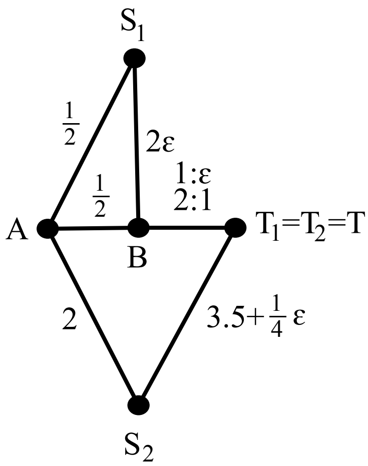

Fig. 23 2-player $AAR_{2\,\text{player}}$ game for *MvP*

Let us now prove $MvP^{\text{ISD}} = \infty$. We will use the game in Figure 23, which is a slightly modified version of the game in Figure 22. Again, this construction causes commitment to a pure strategy to hardly help player 1 at all. player 1's best pure-strategy commitment is still the path $S_1 \to A \to B \to T$, which will be enough to incentivize player 2 to not use the edge $B \to T$, but leaves player 1 with a cost of $1+\varepsilon$. In contrast, if player 1 can commit to a mixed strategy, she can randomize between $S_1 \to A \to B \to T$ ε of the time, and the shortest path ($S_1 \to B \to T$) the rest of the time. This is enough to incentivize player 2 to not use the edge between B and T, causing the expected cost for player 1 to be $\varepsilon(1+\varepsilon)+(1-\varepsilon)(2\varepsilon) = 3\varepsilon - \varepsilon^2$. Letting $\varepsilon \to 0$, we obtain $MvP^{\text{ISD}} = \infty$. ■

6 Experimental Results

Our theoretical results up to this point demonstrate that in many classes of games, there exist games with high values of commitment. However, this is merely a bound on how high these values might grow for a given class of games and does not say anything about what the expected value of commitment might be or show how typical such games are. It is not unreasonable to expect that even when a class has a instance with an unbounded value of commitment there might exist subclasses where at most a small number of outliers have a high value of commitment, but most games in the subclass have much smaller values. Below we focus on commonly studied subclasses of normal-form games, and show experimentally that there do exist subclasses where commitment is valuable and others where it adds no value to the leader. For this experimental evaluation, we randomly generated two-player games using the GAMUT software package [18], which is the standard package to use for such experiments.

We generated random games from 16 families: bidirectional local-effect games (LEG), covariant games, dispersion games, grab the dollar, guess two thirds of the average, location games, majority voting, minimum effort games, polymatrix games, random games, random graphical games, random local-effect games, random zero-sum, traveler's dilemma, uniform LEG, and war of attrition. For each game family, we kept the number of actions available to player 1 equal to the number of actions available to player 2 and we varied the normal-form game size from 5×5 to 40×40 with step size 5. Note that some of the game families that we consider here may have versions with continuous action spaces. In the GAMUT package, the action spaces of such games are always discretized. For example, in grab the dollar games, rather than allowing each player to grab the prize at any time, each player is presented with a set of times at which she can attempt to grab the prize. In each game, the utilities are normalized to be from 0 to 1. We left any other parameters at their default values. For each game size and game family, we generated 50 game instances. For each game instance, we computed the following solutions:

1. an optimal mixed strategy to commit to for player 1,
2. an optimal pure strategy to commit to for player 1,
3. a Nash equilibrium with maximum utility for player 1,
4. a Nash equilibrium with minimum utility for player 1,
5. a Nash equilibrium with maximum social welfare,
6. a correlated equilibrium with maximum utility for player 1,
7. a correlated equilibrium with minimum utility for player 1,
8. a correlated equilibrium with maximum social welfare.

The first solution was computed using an LP [5], the second using a brute force search, 3-5 were computed using a MIP [25], and 6-8 using the standard correlated equilibrium LP. The plots in Figures 24–39 show the utility for player 1 for each game family, game size, and solution concept, averaged over 50 game instances. We used the CPLEX 12.4 solver for the MIPs and LPs. For each game family, we split the eight plots into three subfigures:

(a) the utility for player 1 from committing to an optimal mixed (solid red line) or pure (dashed black line) strategy;
(b) the maximum (solid red line) and minimum (dotted blue line) Nash equilibrium utilities for player 1, and the utility for player 1 in a welfare-maximizing NE (dashed black line);
(c) the maximum (solid red line) and minimum (dotted blue line) correlated equilibrium utilities for player 1, and the utility for player 1 in a welfare-maximizing CE (dashed black line).

The family of random games has a special place among the 16 families. It includes all two-player normal-form simultaneous-move games with finite action sets for both players, with the utilities drawn from a uniform distribution. Each of the other 15 game families imposes its own restrictions on the player's utilities. Thus, the family of random games is a superset of all 16 families. We would like to use the experimental results for the random games as an example to demonstrate how the questions that we want to answer in this section are different from the theoretical results obtained in the earlier sections: while it follows from Corollary 1 that the value of both pure and mixed commitment is infinite for this family of games, we can see (Figure 33) that both the expected utility of commitment for the leader and the maximum NE utility for player 1 approach 1 as the game size approaches infinity. Thus, the expected value of both pure and mixed commitment is 1 for random games. As another example of how the expected value of commitment is different from the maximum value of commitment, consider random zero-sum games (Figure 36). It follows from Theorem 1 that the value of both pure and mixed commitment is equal to 1 in constant-sum games. However, for large random zero-sum games, the leader is likely to get only a very small utility from committing to any pure strategy, and thus the expected value of pure commitment approaches 0 as the game size approaches infinity. We will discuss the experimental results for this and other game families in more detail below.

There are two game families in which all the eight lines in the corresponding subfigures (a)-(c) coincide: guess two thirds of the average games and location games. In these games, the optimal mixed strategy to commit to for player 1 is a pure strategy which is also part of a pure-strategy Nash equilibrium of the game. Moreover, both players' utilities are constant across all CE of each game. We will explain why this happens in guess two thirds of the average. In this game, each player chooses a number from the given set of n numbers. If the two players choose different numbers, the player whose number is closer to the $2/3$ of the average of the two numbers gets a utility of 1 and the other player gets 0. If the players choose the same number, each gets a utility of 0.5. In the unique NE of this game, both players choose the smallest of the given numbers, which results in a utility of 0.5 for each player. The same strategy profile is also the unique CE of the game.

In three game families, all the plots coincide except the minimum utility for player 1 in NE and CE: dispersion games, minimum effort games, random LEG. Similarly to the game families discusses in the previous paragraph, the utility for player 1 in Stackelberg games (with commitment to either pure or mixed strategies) is the same as in leader-utility-maximizing NE or CE. The difference is that there exists an NE with a low utility for player 1. Consider the minimum effort games. In this game, each player chooses a level of effort from a given set. The player's utility is equal to the minimum effort level among the two players minus a fraction of the player's own effort level. There is a set of pure NE in this game in which both players choose the same effort level. Both players' utilities are the highest when both players choose the highest effort level. The leader can guarantee herself the highest utility by committing to the highest effort level, which is also the strategy played by both players in the welfare-maximizing and leader-utility-maximizing NE and CE. However, since there exists an NE in which both players get a low utility (namely, the NE in which both players choose a low effort level), the dotted lines in subfigures (b) and (c) stand out from the solid and dashed lines in subfigures (a)-(c) of Figure 31.

We will now discuss the remaining 11 game families.

Commitment vs. Nash equilibrium. Since a Stackelberg strategy maximizes the leader's utility, we will compare the utility for player 1 in the NE which maximizes player 1's utility to the leader's utility of committing to a mixed strategy. In the plots, that corresponds to comparing the solid red line in each subfigure (a) to the solid red line in the corresponding subfigure (b). We can see that the leader can benefit from commitment in covariant games, majority voting, polymatrix games, random games, random graphical games, and traveler's dilemma. The value of commitment is especially high in traveler's dilemma games. Both players get utilities close to $1/n$ in the unique NE in these games, while the leader's utility from committing to a mixed strategy is close to 1 in $n \times n$ games. In this variant of traveler's dilemma, each player can choose from one of n price levels, and the player who declares the lower price gets a utility equal to that price plus a small reward. The Stackelberg strategy is to declare the highest price, while in NE both players declare the lowest price. Thus, the value of commitment approaches infinity in such games.[3]

Even in cases when the leader's highest NE utility is equal to the leader's Stackelberg utility, commitment may still be valuable. In such games, the ability to commit lets the leader choose the NE with the highest utility for the leader instead of the other NE with lower utilities. We can see that the leader can benefit from choosing the NE in all games except guess two thirds of the average, location games, and random zero-sum.

Commitment vs. correlated equilibrium. Next, we compare the utility for player 1 in the CE which maximizes player 1's utility to the leader's utility of playing a Stackelberg strategy. This corresponds to comparing the solid red lines in subfigures (a) and (c). Since every NE is also a CE, there are fewer game families in which the leader clearly benefits from commitment than we found when comparing Stackelberg and NE: covariant games, majority voting, polymatrix games, random games, random graphical games, and traveler's dilemma. Note that a strictly dominated strategy is never played in a correlated equilibrium. Thus, in order to construct an example in which player 1 gets a higher utility from commitment than she gets in any CE, it is enough to construct a game in which the optimal strategy to commit to is dominated.

[3]Due to numerical issues, our MIP-based Nash solver computes a higher NE utility for player 1 in some traveler's dilemma games. This happens because there is an epsilon-NE with extremely small epsilon in which each player randomizes over several strategies. We think this also says something interesting about this game, namely that extremely small deviations from rationality are sufficient to completely change the game-theoretic prediction and thereby the point this game is intended to make.

Commitment to mixed strategies vs. commitment to pure strategies. This corresponds to comparing the two lines in each subfigure (a). We can see that there are a number of games in which the leader gets the same utility from committing to a pure strategy as she gets from committing to a mixed strategy. For example, consider grab the dollar games. Each player's action set is the set of times when the player can grab the prize. If the two players both grab the prize at the same time, it will be torn, and each player will get a utility of 0. Otherwise, the first player to grab the dollar gets to keep it. Player 1's optimal mixed strategy to commit to is to declare that she is going to grab the dollar at the earliest opportunity. The follower will then choose a later time, and the leader get the utility of 1. Since the optimal mixed strategy is actually a pure strategy, the two lines coincide.

The biggest difference between committing to a mixed strategy vs. committing to a pure strategy is in random zero-sum games. In this games family, the two lines diverge: as the game size grows, the leader's utility from committing to a mixed strategy approaches 0.5, while the leader's utility from committing to a pure strategy approaches 0. We can provide an intuition as to why this happens as follows. In zero-sum games, every Stackelberg strategy is also an NE strategy, and each player gets the same utility across all NE profiles. Since the distribution from which we generate random zero-sum games is symmetric with respect to the two players, the expected utility that player 1 gets in an NE must be equal to the expected utility that player 2 gets in an NE. The sum of the expected utilities that the players get in NE is 1, thus each player's expected utility in an NE is 0.5, which is also the expected utility from committing to a Stackelberg strategy for any player. In the pure commitment case, the leader's utility approaches 0 as n grows. We can show that for any fixed number x, the expected number of rows to which the leader can commit and get the utility of at least x approaches 0 as n grows. We will think of each row as a point in $[0,1]^n$. The subspace of rows $S_x \subset [0,1]^n$ in which every element is greater than or equal to x has volume $(1-x)^n$. As n grows, the volume of S_x shrinks exponentially while the number of rows available to the leader to choose from grows only linearly. Thus the expected number of leader's strategies in S_x approaches 0 as n grows, which implies that the probability of there being a row with all elements greater than or equal to x also approaches 0.

While we performed our experimental results on GAMUT games because this is the standard for the field, GAMUT does not provide game generators for all of the classes of games that we studied in the theoretical part of the paper. Future research could be devoted to creating new generators for these classes of games and extending our experimental results to them.

7 Conclusion

On the one hand, the notion of the value of commitment studied in this paper fits well among various other notions that aim to quantify strategic effects in games, including the price of anarchy, the price of stability, the value of mediation, etc. On the other hand, to our knowledge, this concept is unique among these concepts in the sense that it focuses on the benefit of a strategic tool—commitment—to a specific player, rather than the cost of strategic behavior to welfare in general. (It should be noted that the value of mediation also considers the benefit of a strategic tool—correlated strategies—to social welfare, though this is a somewhat different type of strategic tool, one that is used by the players collectively rather than by an individual player.) Additionally, we evaluated the value of commitment experimentally. While we performed our experimental results on GAMUT games because this is the standard for the field, GAMUT does not provide game generators for all of the classes of games that we studied in the theoretical part of the paper. Future research could be devoted to creating new generators for these classes of games and extending our experimental results to them. One additional obvious direction for future research concerns what happens in symmetric routing games with more than 2 players. We believe that there are many other directions for future research, including investigating the value of commitment more generally, as well as investigating the value of other strategic tools.

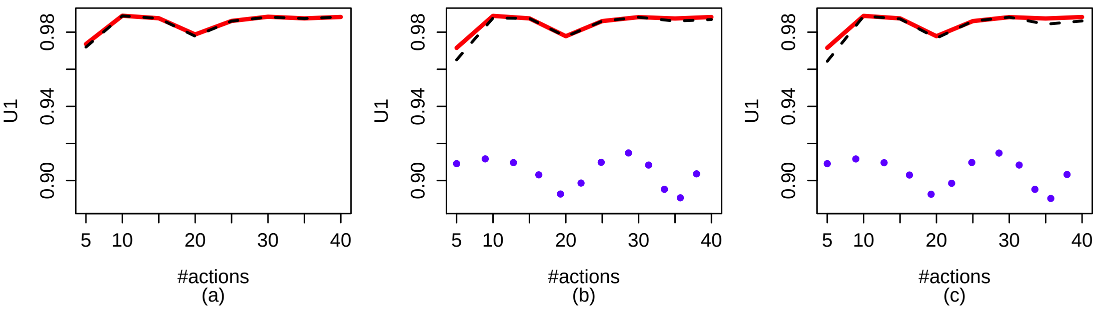

Fig. 24 Bidirectional local-effect games (LEG).

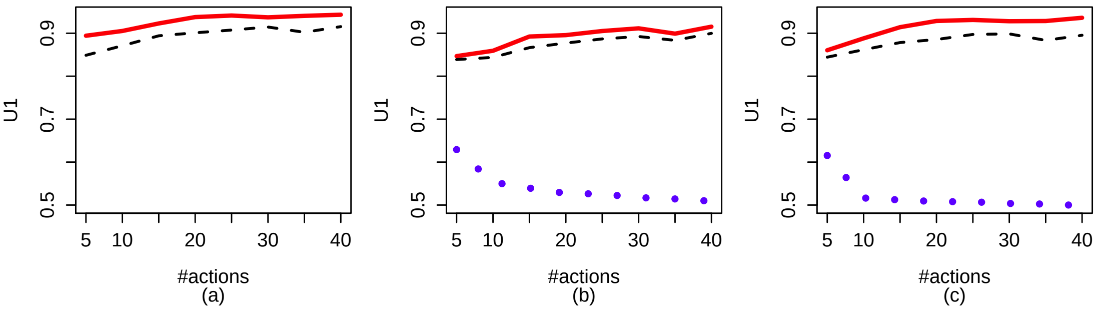

Fig. 25 Covariant games.

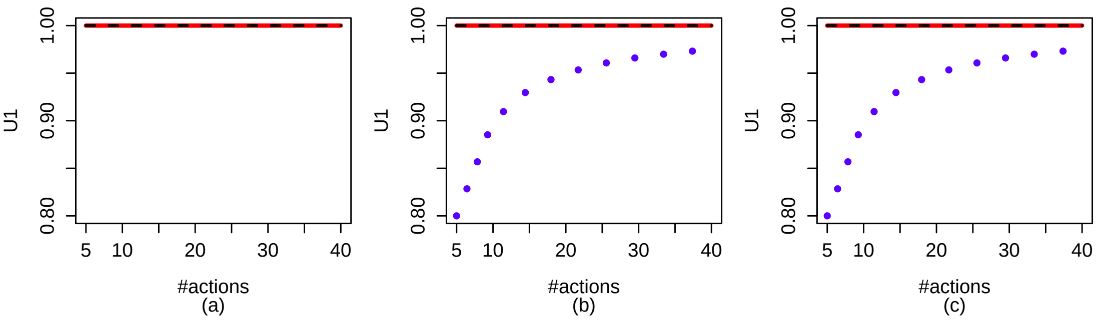

Fig. 26 Dispersion games.

nn

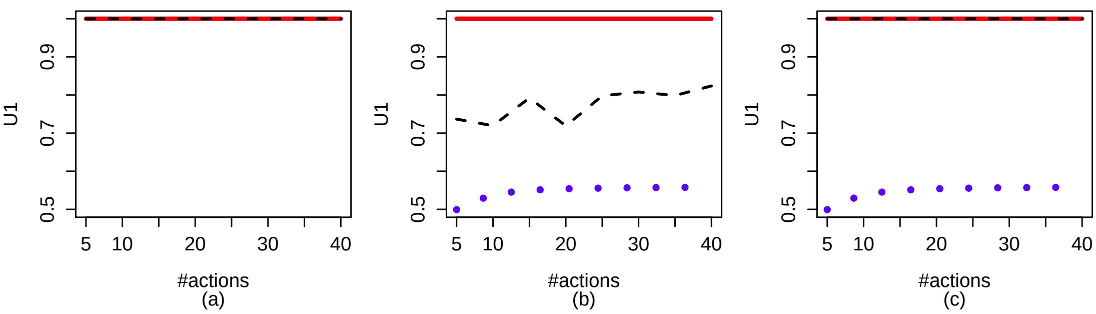

Fig. 27 Grab the dollar games.

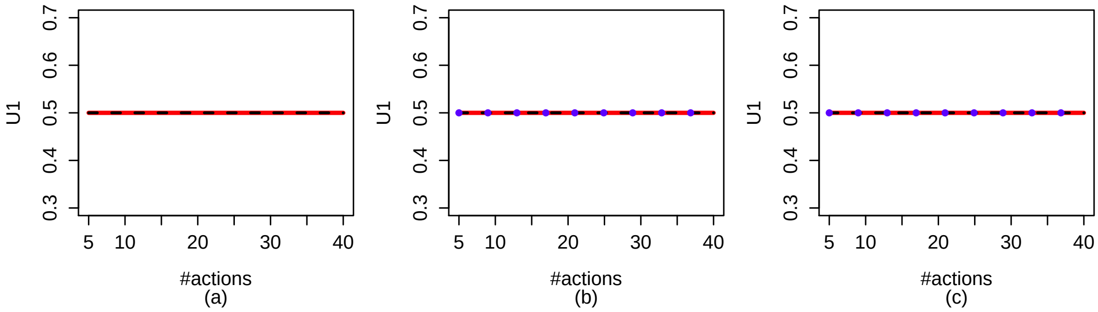

Fig. 28 Guess two thirds of the average games.

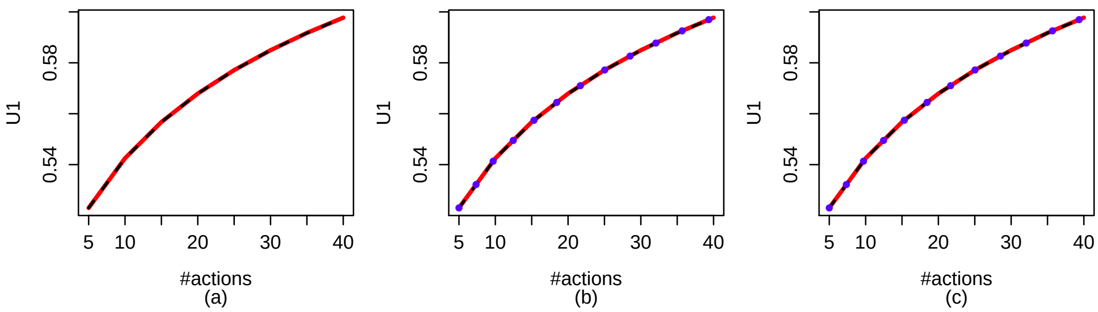

Fig. 29 Location games.

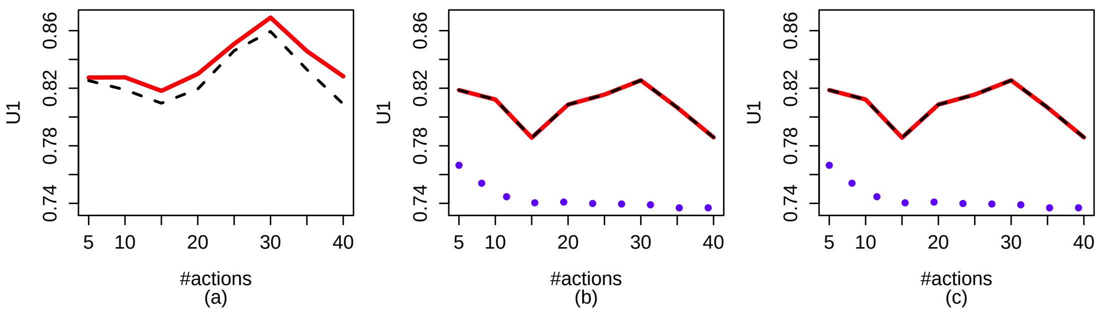

Fig. 30 Majority voting games.

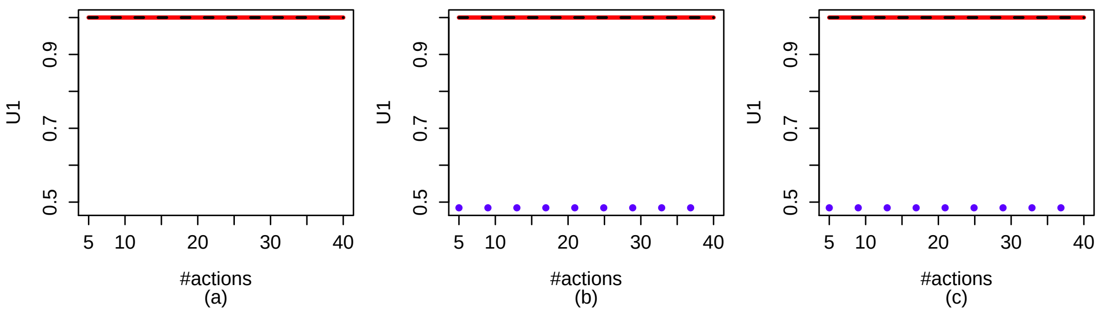

Fig. 31 Minimum effort games.

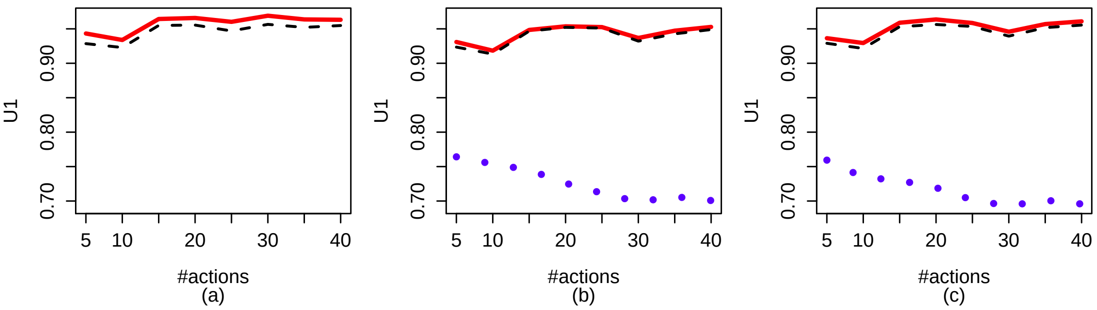

Fig. 32 Polymatrix games.

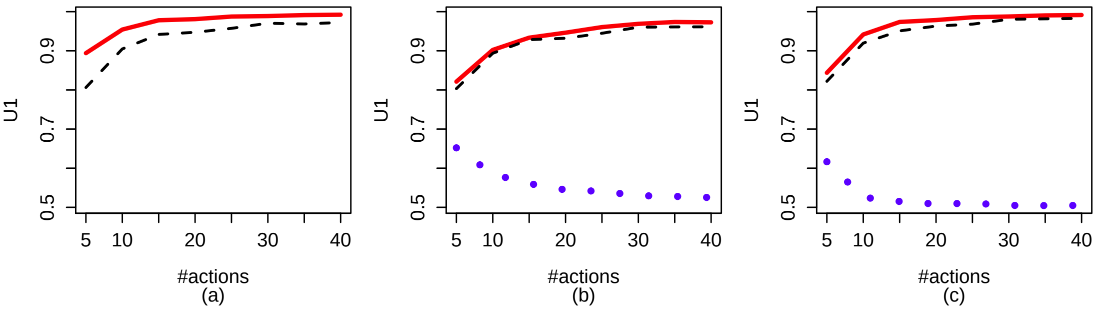

Fig. 33 Random games.

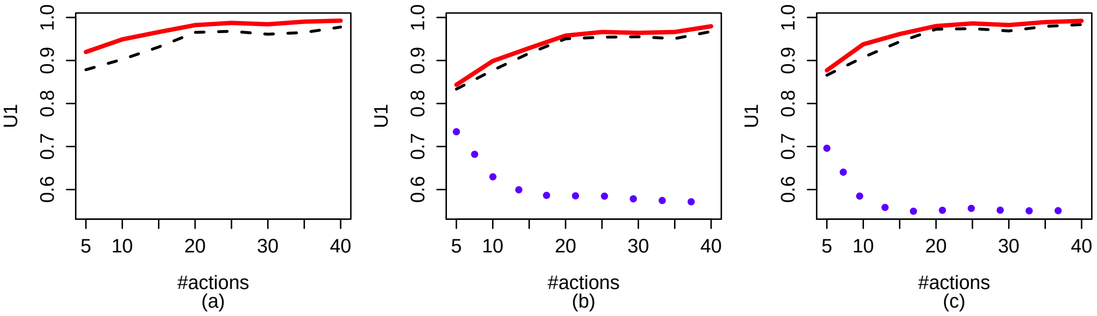

Fig. 34 Random graphical games.

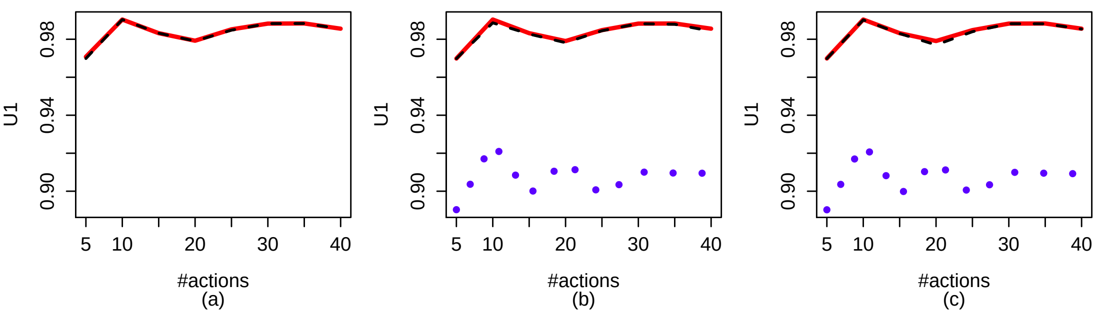

Fig. 35 Random LEG.

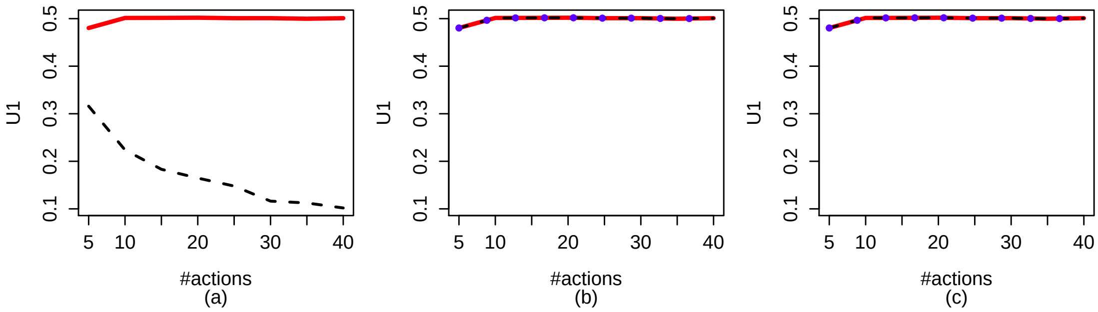

Fig. 36 Random zero-sum games.

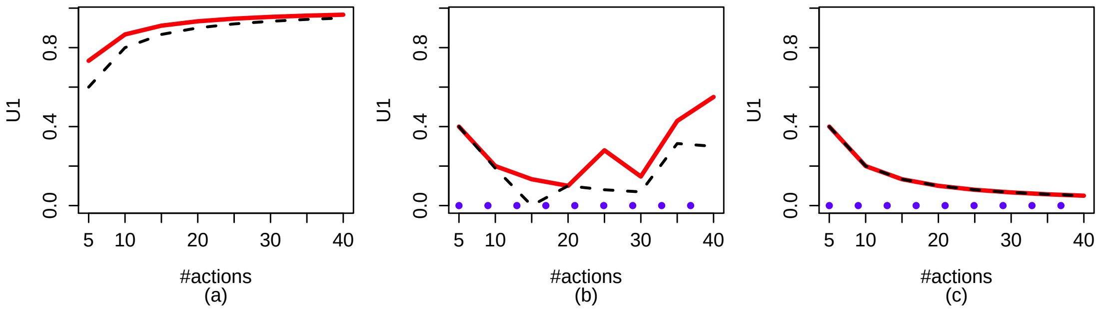

Fig. 37 Traveler's dilemma games.

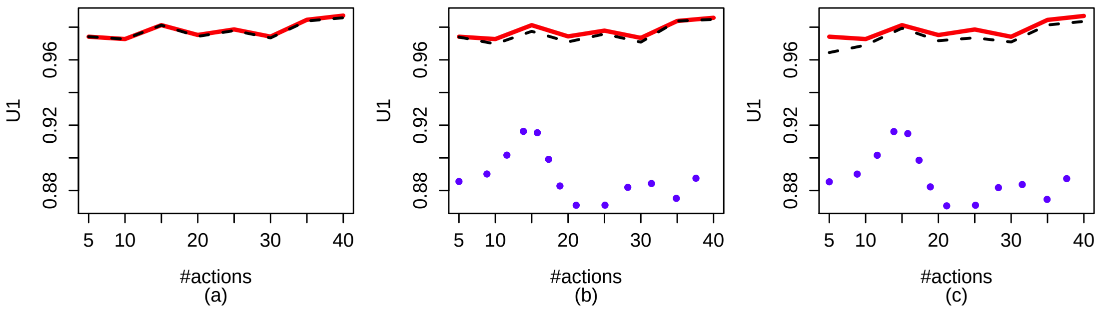

Fig. 38 Uniform Local-Effect games.

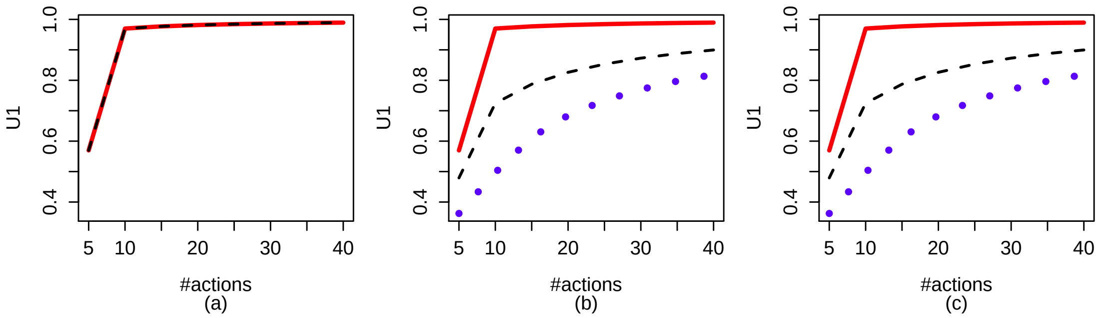

Fig. 39 War of attrition games.

Acknowledgments

This work has been supported by NSF Awards IIS-0812113, IIS-0953756, CCF-1101659, ARO Grants W911NF-09-1-0459, W911NF-11-1-0332, W911NF-12-1-0550, and an Alfred P. Sloan fellowship. We would like to thank Kamesh Munagala for his suggestion to consider games that are close to constant-sum.

References

1. Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 295–304, 2004.
2. Itai Ashlagi, Dov Monderer, and Moshe Tennenholtz. On the value of correlation. *Journal of Artificial Intelligence Research*, 33:575–613, 2008.
3. Vincenzo Bonifaci, Tobias Harks, and Guido Schäfer. Stackelberg routing in arbitrary networks. *Math. Oper. Res.*, 35(2):330–346, May 2010.
4. George Christodoulou and Elias Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Proceedings of the 13th Annual European Symposium*, pages 59–70, 2005.
5. Vincent Conitzer and Dmytro Korzhyk. Commitment to correlated strategies. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 632–637, San Francisco, CA, USA, 2011.
6. Vincent Conitzer and Tuomas Sandholm. Computing the optimal strategy to commit to. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 82–90, Ann Arbor, MI, USA, 2006.
7. Antoine Augustin Cournot. *Recherches sur les principes mathématiques de la théorie des richesses (Researches into the Mathematical Principles of the Theory of Wealth)*. Hachette, Paris, 1838.
8. Dimitris Fotakis. Stackelberg strategies for atomic congestion games. In Lars Arge, Michael Hoffmann, and Emo Welzl, editors, *Algorithms ? ESA 2007*, volume 4698 of *Lecture Notes in Computer Science*, pages 299–310. Springer Berlin Heidelberg, 2007.
9. Manish Jain, James Pita, Milind Tambe, Fernando Ordóñez, Praveen Paruchuri, and Sarit Kraus. Bayesian Stackelberg games and their application for security at Los Angeles International Airport. *SIGecom Exch.*, 7(2):1–3, 2008.
10. Christopher Kiekintveld, Manish Jain, Jason Tsai, James Pita, Fernando Ordóñez, and Milind Tambe. 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)*, pages 689–696, Budapest, Hungary, 2009.
11. Yannis A. Korilis, Aurel A. Lazar, and Ariel Orda. Achieving network optima using Stackelberg routing strategies. *IEEE/ACM Transactions on Networking*, 5(1):161–173, 1997.
12. Dmytro Korzhyk, Vincent Conitzer, and Ronald Parr. Complexity of computing optimal Stackelberg strategies in security resource allocation games. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 805–810, Atlanta, GA, USA, 2010.
13. Dmytro Korzhyk, Zhengyu Yin, Christopher Kiekintveld, Vincent Conitzer, and Milind Tambe. Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness. *Journal of Artificial Intelligence Research*, 41(2):297–327, 2011.
14. Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In *Symposium on Theoretical Aspects in Computer Science*, pages 404–413, 1999.
15. Joshua Letchford and Vincent Conitzer. Computing optimal strategies to commit to in extensive-form games. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 83–92, Cambridge, MA, USA, 2010.
16. Joshua Letchford, Vincent Conitzer, and Kamesh Munagala. Learning and approximating the optimal strategy to commit to. In *Proceedings of the Second Symposium on Algorithmic Game Theory (SAGT-09)*, pages 250–262, Paphos, Cyprus, 2009.
17. Joshua Letchford, Liam MacDermed, Vincent Conitzer, Ronald Parr, and Charles Isbell. Computing optimal strategies to commit to in stochastic games. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 1380–1386, Toronto, ON, Canada, 2012.
18. Eugene Nudelman, Jennifer Wortman, Kevin Leyton-Brown, and Yoav Shoham. Run the GAMUT: A comprehensive approach to evaluating game-theoretic algorithms. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 880–887, New York, NY, USA, 2004.
19. Christos H. Papadimitriou. Algorithms, games and the Internet. In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*, pages 749–753, 2001.
20. Praveen Paruchuri, Jonathan P. Pearce, Janusz Marecki, Milind Tambe, Fernando Ordóñez, and Sarit Kraus. Playing games for security: An efficient exact algorithm for solving Bayesian Stackelberg games. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 895–902, Estoril, Portugal, 2008.
21. James Pita, Manish Jain, Fernando Ordóñez, Christopher Portway, Milind Tambe, and Craig Western. Using game theory for Los Angeles airport security. *AI Magazine*, 30(1):43–57, 2009.
22. James Pita, Milind Tambe, Chris Kiekintveld, Shane Cullen, and Erin Steigerwald. GUARDS - Game theoretic security allocation on a national scale. In *Proceedings of the Tenth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 37–44, Taipei, Taiwan, 2011.
23. Tim Roughgarden. Stackelberg scheduling strategies. *SIAM Journal on Computing*, 33(2):332–350, 2004.
24. Tim Roughgarden and Éva Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
25. Tuomas Sandholm, Andrew Gilpin, and Vincent 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.
26. Eric Shieh, Bo An, Rong Yang, Milind Tambe, Craig Baldwin, Joseph DiRenzo, Ben Maule, and Garrett Meyer. PROTECT: A deployed game theoretic system to protect the ports of the United States. In *Proceedings of the Eleventh International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Valencia, Spain, 2012.

27. Jason Tsai, Shyamsunder Rathi, Christopher Kiekintveld, Fernando Ordonez, and Milind Tambe. IRIS - a tool for strategic security allocation in transportation networks. In *Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 37–44, Budapest, Hungary, 2009.
28. Heinrich von Stackelberg. *Marktform und Gleichgewicht*, pages 58–70. Springer, Vienna, 1934.
29. Bernhard von Stengel and Shmuel Zamir. Leadership games with convex strategy sets. *Games and Economic Behavior*, 69:446–457, 2010.