Approximate Strategic Reasoning through Hierarchical Reduction of Large Symmetric Games

Michael P. Wellman, Daniel M. Reeves, Kevin M. Lochner, Shih-Fen Cheng, and **Rahul Suri**
University of Michigan
Ann Arbor, MI 48109-2110 USA
{wellman,dreeves,klochner,chengsf,rsuri}@umich.edu

Abstract

To deal with exponential growth in the size of a game with the number of agents, we propose an approximation based on a hierarchy of reduced games. The reduced game achieves savings by restricting the number of agents playing any strategy to fixed multiples. We validate the idea through experiments on randomly generated local-effect games. An extended application to strategic reasoning about a complex trading scenario motivates the approach, and demonstrates methods for game-theoretic reasoning over incompletely-specified games at multiple levels of granularity.

Motivation

Consider the task of selecting among a large set of strategies to play in an 8-player game. Through careful judgment you manage to narrow down the candidates to a reasonable number of strategies (say 35). Because the performance of a strategy for one agent depends on the strategies of the other seven, you wish to undertake a game-theoretic analysis of the situation. Determining the payoff for a particular strategy profile is expensive, however, as your observations of prior game instances are quite limited, and the only *operational* description of the game is in the form of a simulator that takes a non-negligible time (say 10 minutes) to produce one outcome. Moreover, since the environment is stochastic, numerous samples (say 12) are required to produce a reliable estimate for even one profile. At two hours per profile, exhaustively exploring profile space will require $2 \cdot 35^8$ or 4.5 trillion hours simply to estimate the payoff function representing the game under analysis. If the game is symmetric, you can exploit that fact to reduce the number of distinct profiles to $\binom{42}{8}$, which will require 236 million hours. That is quite a bit less, but still much more time than you have.

This is the situation we face as entrants in the annual Trading Agent Competition (TAC) travel-shopping market game (Wellman *et al.* 2003). The necessity of empirical evaluation in this setting combined with the infeasibility of exhaustive analysis prompts us to seek principled ways to direct a non-exhaustive exploration. In this paper we investigate the exploitation of hierarchical structure in the space of profiles to balance the goals of spanning the overall space and focusing effort on the most promising regions.

Copyright © 2005, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

The idea is that although a strategy's payoff does depend on the play of other agents (otherwise we are not in a game situation at all), it may be relatively insensitive to the exact numbers of other agents playing particular strategies. For example, let $(s, m; s')$ denote a profile where m other agents play strategy s, and the rest play s'. In many natural games, the payoff for playing any particular strategy against this profile will vary smoothly with m. If such is the case, we sacrifice relatively little fidelity by restricting attention to subsets of profiles, for instance those with only even numbers of any particular strategy. To do so essentially transforms the N-player game to an $N/2$-player game over the same strategy set, where the payoffs to a profile in the reduced game are simply those from the original game where each strategy in the reduced profile is played twice.

The potential savings from reduced games are considerable, as they contain combinatorially fewer profiles. The 4-player approximation to the TAC game (with 35 strategies) comprises 73,815 distinct profiles, compared with 118 million for the original 8-player game. In case exhaustive consideration of the 4-player game is still infeasible, we can approximate further by a corresponding 2-player game, which has only 630 profiles. Approximating by a 1-player game is tantamount to ignoring strategic effects, considering only the 35 "profiles" where the strategies are played against themselves. In general, an N-player symmetric game with S strategies includes $\binom{N+S-1}{N}$ distinct profiles. Figure 1

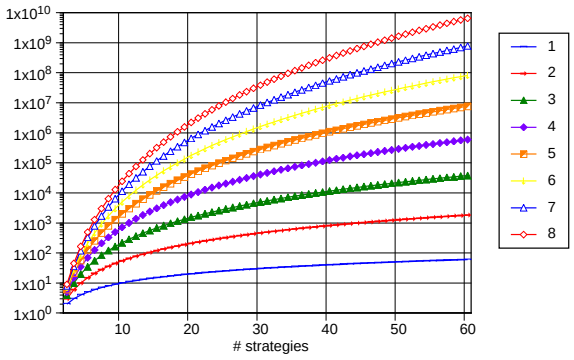


Figure 1: Number of distinct profiles (log scale) of a symmetric game, for various numbers of players and strategies.

shows the exponential growth in both N and S.

The main contribution of this paper is to introduce the concept of approximation through hierarchical game reduction. We begin by providing a more precise definition of reduced games. Next we present evidence supporting the approximation of games by reduced versions, based on theoretical and experimental studies of three game classes. The rest of the paper focuses on our application to the TAC travel-shopping domain, demonstrating the use of game-theoretic reasoning about possible equilibria based on only incomplete evaluation of the underlying payoff function.

Hierarchy of Reduced Games

We develop our hierarchical reduction concepts in the framework of *symmetric normal-form games*.[1]

Definition 1 $\Gamma = \langle N, \{S_i\}, \{u_i()\}\rangle$ *is an* N-player normal-form game, *with* strategy set S_i *the available strategies for player* i, *and the* payoff function $u_i(s_1, \ldots, s_N)$ *giving the utility accruing to player* i *when players choose the* strategy profile $(s_1, \ldots, s_N)$.

Definition 2 *A normal-form game is* symmetric *if the players have identical strategy spaces* ($S_i = S$) *and* $u_i(s_i, s_{-i}) = u_j(s_j, s_{-j})$, *for* $s_i = s_j$ *and* $s_{-i} = s_{-j}$ *for all* $i, j \in \{1, \ldots, N\}$. *Thus we can denote a symmetric game by* $\langle N, S, u()\rangle$, *with* $u(t, s)$ *the payoff to any player playing strategy* t *when the remaining players play profile* s.

Our central concept is that of a *reduced* game.

Definition 3 *Let* $\Gamma = \langle N, S, u()\rangle$ *be an* N*-player symmetric game, with* $N = pq$ *for integers* p *and* q. *The* p-player reduced version of Γ, *written* $\Gamma\downarrow_p$, *is given by* $\langle p, S, \hat{u}()\rangle$, *where*

$$\hat{u}_i(s_1, \ldots, s_p) = u_{q \cdot i}(\underbrace{s_1, \ldots}_{q}, \underbrace{s_2, \ldots}_{q}, \ldots, \underbrace{s_p, \ldots}_{q}).$$

In other words, the payoff function in the reduced game is obtained by playing the specified profile in the original q *times.*

The idea of a reduced game is to coarsen the profile space by restricting the degrees of strategic freedom. Although the original set of strategies remains available, the number of agents playing any strategy must be a multiple of q. Every profile in the reduced game is one in the original game, of course, and any profile in the original game can be reached from a profile contained in the reduced game by changing at most $p(q-1)$ agent strategies.

To search a profile space hierarchically, we apply a series of game reductions. The game resulting from such a series is independent of the reduction ordering. Let $q = r \cdot r'$. Then

$$(\Gamma\downarrow_{p \cdot r})\downarrow_p = (\Gamma\downarrow_{p \cdot r'})\downarrow_p = \Gamma\downarrow_p .$$

Example 1 (FPSBn) *In the* n*-player first-price sealed-bid auction, player* i *has a private value* v_i, *decides to bid* b_i, *and obtains payoff* $v_i - b_i$ *if its bid is highest (and zero otherwise). We define* FPSBn *as a special case where* $v_i \sim U[0, 1]$, *and agents are restricted to parametrized strategies, bidding* $k_i v_i$ *for* $k_i \in [0, 1]$.

Let $n = pq$. In the reduced game $\mathrm{FPSB}n\downarrow_p$, each agent $i = 1, \ldots, p$ selects a single action k_i, which then gets applied to q valuations $v_{i_1}, \ldots, v_{i_q}$ to define q bids. The auction proceeds as normal, and agent i's payoff is defined as the *average* payoff associated with its q bids. Note that the game $\mathrm{FPSB}n\downarrow_p$ is quite a different game from either FPSBn or FPSBp. When represented explicitly over a discrete set of actions, $\mathrm{FPSB}n\downarrow_p$ is the same size as FPSBp, and both are exponentially smaller than FPSBn.

Reduced-Game Approximations

Our premise is that the reduced game will often serve as a good approximation of the full game it abstracts. We know that in the worst case it does not. In general, an equilibrium of the reduced game may be arbitrarily far from equilibrium with respect to the full game, and an equilibrium of the full game may not have any near neighbors in the reduced game that are close to equilibrium there.[2] The question, then, is whether useful hierarchical structure is present in "typical" or "natural" games, however we might identify such a class of games of interest. Although we have no general characterization of the class of games for which the approximation works well, we provide positive evidence for three specific classes of symmetric games in the sections below.

FPSBn

The n-player FPSB auction has a unique symmetric Nash equilibrium, at $k = \frac{n-1}{n}$ (Krishna 2002). For example, the equilibrium for FPSB2 is $1/2$, and for FPSB4 it is $3/4$. From the following theorem, giving the equilibrium of $\mathrm{FPSB}n\downarrow_p$, we have $2/3$ in equilibrium for $\mathrm{FPSB}4\downarrow_2$.

Theorem 1 *The unique symmetric Nash equilibrium of* $\mathrm{FPSB}n\downarrow_p$ *is*

$$\frac{n(p-1)}{p + n(p-1)}.$$

(Proofs of this and subsequent theorems are omitted.)

The difference between equilibrium strategies of $\mathrm{FPSB}4\downarrow_2$ and FPSB4 is one measure of their distance. An alternative measure of approximation quality is to evaluate solutions of the reduced game in the context of the original. Specifically, we ask: If the agents play a reduced-game equilibrium in the original game, how much can a single agent gain from deviating from such a profile? If the answer is zero, then the equilibria coincide. More generally, the smaller the gain from deviating, the more faithful the reduced game approximation.

[1] Although the methods may generalize to some degree to partially symmetric games, or to exploit extensive forms, we do not pursue such extensions here.

[2] $\mathrm{FPSB}n\downarrow_1$ is an example (albeit a degenerate one) of a reduced game having very different equilibria than the full game. The optimal strategy in the 1-player reduction is to bid zero (as it is in FPSB1) whereas in the n-player game, for $n > 1$, equilibrium players bid a large fraction (at least $1/2$) of their types.

Let us denote by $\epsilon_\Gamma(s)$ the potential gain to deviating from strategy profile s in game Γ. For symmetric game $\Gamma = \langle N, S, u() \rangle$,

$$\epsilon_\Gamma(s) = \max_{s' \in S} u(s', s) - u(s, s). \quad (1)$$

This usage follows the standard notion of approximate equilibrium. Profile s is an $\epsilon_\Gamma(s)$-Nash equilibrium of Γ, with 0-Nash corresponding to exact equilibrium. Henceforth, we drop the game subscript when understood in context.

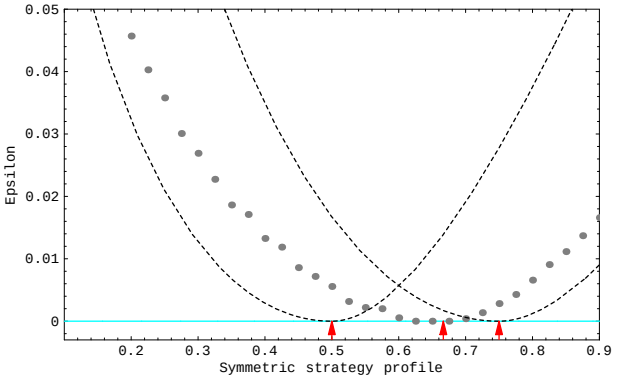


Figure 2: Epsilons for symmetric profiles of FPSB2 (left dashed curve), FPSB4$\downarrow_2$ (dots), and FPSB4 (right dashed curve). Unique symmetric equilibria $(1/2, 2/3, 3/4)$ are indicated by arrows on the x-axis.

Figure 2 plots $\epsilon(k)$ for the three game variations. We derived a closed-form expression for $\epsilon_{\text{FPSB}n}$, whereas the curve for $\epsilon_{\text{FPSB4}\downarrow_2}$ was estimated numerically (though we know its exact root by Theorem 1). Our estimation procedure considered all profiles over discrete values of k, at intervals of $1/40$. At this granularity, FPSB4 comprises 158 times as many profiles as does FPSB4$\downarrow_2$. Based on our analysis, FPSB4$\downarrow_2$ compares quite favorably to FPSB2 as an approximation of FPSB4. In particular, taking their respective equilibrium values, $\epsilon_{\text{FPSB4}}(2/3)$ is nearly ten times smaller than $\epsilon_{\text{FPSB4}}(1/2)$.

We can generalize this conclusion to arbitrary n and p. Let $s^*(\Gamma)$ denote the unique symmetric equilibrium for Γ, which is well-defined for our FPSB games of interest.

Theorem 2 *For all* $n > p > 1$,

$$\epsilon_{\text{FPSB}n}(s^*(\text{FPSB}n\downarrow_p)) < \epsilon_{\text{FPSB}n}(s^*(\text{FPSB}p)).$$

We can also confirm that for any number of players, less drastic FPSB reductions provide better approximations.

Theorem 3 *For all* $n > p > q \geq 1$,

$$\epsilon_{\text{FPSB}n}(s^*(\text{FPSB}n\downarrow_p)) < \epsilon_{\text{FPSB}n}(s^*(\text{FPSB}n\downarrow_q)).$$

Bertrand Oligopoly

The preceding analysis is reassuring, but of course we do not actually need to approximate FPSBn, since its general solution is known. To further evaluate the quality of reduced-game approximations, we turn to other natural games of potential interest. Facilitating such studies was precisely the motivation of the authors of GAMUT (Nudelman *et al.* 2004), a flexible software tool for generating random games from a wide variety of well-defined game classes. Using GAMUT, we can obtain random instances of some class, and examine the relation of the original games to versions reduced to varying degrees. The advantage of a generator such as GAMUT is that we can obtain a full game specification quickly (unlike for TAC), of specified size based on our computational capacity for analysis. Moreover, we can sample many instances within a class, and develop a statistical profile of the properties of interest.

The first class of games we examined using GAMUT is a discrete symmetric version of *Bertrand Oligopoly* (BO), which models price competition among a set of identical producers (Mas-Colell, Whinston, & Green 1995, Section 12.C). In the Bertrand model, each firm declares a price, and total demand at the lowest price is divided among those firms offering that price. The payoffs are given by net profit, which is zero for firms priced above the minimum. To generate a BO game in GAMUT, we specify the numbers of agents and discrete price levels (i.e., the actions), and parameters for the demand and cost functions. By specifying ranges for some function parameters, we define a class of such games.

We generated three random 8-player BO games, with 3, 4, and 5 actions, respectively. We measured ϵ_{BO} for each pure profile in the original and reduced games. Figure 3 presents the correlation coefficients, for each BO game instance, of the ϵ_{BO} for profiles shared by each pair of game versions. As we see, the correlation is quite strong between the original and reduced games, with relative values depending on the degree of reduction.

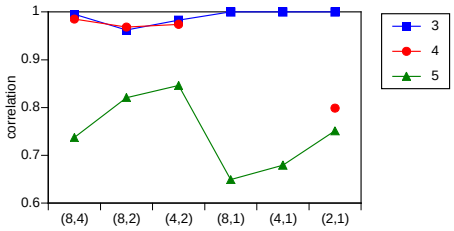


Figure 3: Correlation among ϵ_{BO} values for profiles in BO (BO$\downarrow_8$) and reduced versions, for 3-, 4- and 5-action random BO instances. (q, r) marks the column comparing BO$\downarrow_q$ with BO$\downarrow_r$. The 4-action BO$\downarrow_8$ and BO$\downarrow_4$ showed no variation in ϵ_{BO} for profiles shared with BO$\downarrow_1$.

Local-Effect Games

We performed a more thorough experimental study of a particular class known as *local-effect games* (LEGs) (Leyton-Brown & Tennenholtz 2003), a localized version of congestion games motivated by problems in AI and computer networks. Specifically, we consider symmetric bi-directional local-effect games randomly generated by GAMUT by creating random graph structures and random polynomial payoff functions decreasing in the number of action-nodes chosen.

In a preliminary experiment, we generated 15 symmetric LEG instances with six players and two strategies, and

payoffs normalized on $[0, 1]$. For each of these we generated the corresponding 3-player reduction. We then fed all 30 of these instances to GAMBIT (McKelvey, McLennan, & Turocy 1992), a general game-solving package, which computed the complete set of Nash equilibria for each. In 11 of the original games, all equilibria are pure, and in these cases the equilibria of the reduced games match exactly. In the remaining four games, GAMBIT identified strictly mixed equilibria. In two of these cases, for every equilibrium in the full game there exists an equilibrium of the reduced game with strategy probabilities within 0.1. In the remaining two games, there are long lists of equilibria in the full game and shorter lists in the corresponding reduced games. In these cases, most but not all of the equilibria in the reduced game are approximations to equilibria in the full.

In broader circumstances, we should not expect to see (nor primarily be concerned with) direct correspondence of equilibria in the original and reduced games. Thus, we evaluate the approximation of a reduced game in terms of the average $\epsilon(s^*)$ in the original game over all its equilibrium profiless^* in the reduced game. Note that to calculate this measure, we need not be able to solve the full game. Since the games under consideration are symmetric, our assessment includes only the symmetric equilibria, where every agent plays the same (mixed) strategy.[3]

We next evaluated 2-strategy local-effect games with n players, for $n \in \{4, 6, 8, 10, 12\}$, generating 200–10,000 random instances of each. Figure 4 shows the average $\epsilon(s^*)$ for every possible reduction of every game, starting with the most drastic reduction—to one player—and ending with the highest-fidelity reduction, i.e., to half as many players. We also include the average ϵ for the social optimum (the profile maximizing aggregate payoff) in each game class as calibration. We find that the social optimum fares better than the equilibria in the 1-player reduction (i.e., the strategy yielding the highest payoff if played by everyone) but that all the higher fidelity reductions yield equilibria with average ϵ significantly better. The only exception is the case of approximating 4-player games with their 2-player reductions. We note that in fully 90% of the 4-player LEG instances, the social optimum is also an equilibrium, making it particularly hard to beat for that game class. The percentages are also high in the other classes—decreasing with the number of players to 77% for the 12-player instances—yet the social optima in all the other classes are beaten by solutions to the reduced games.

In addition to confirming the hypothesis that we can approximate a game by its reduction, we conclude from this analysis that, as in FPSBn, we get diminishing returns on refinement. There is a large benefit to going from 1- to 2-player games (i.e., bringing in strategic interactions at all), then progressively less by adding more fidelity.

Developing TAC Agents

TAC travel-shopping is an 8-player symmetric game, with a complex strategy space and pivotal agent interactions.

[3]Symmetric games necessarily have symmetric equilibria (Nash 1951), though they may have asymmetric equilibria as well.

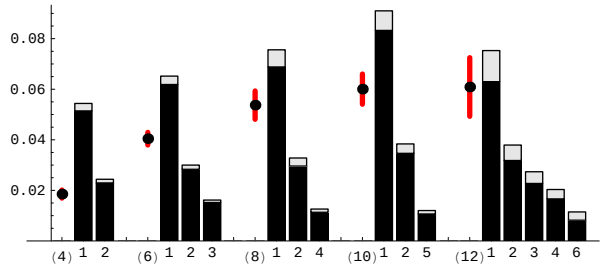


Figure 4: Local-effect games with 4, 6, 8, 10, and 12 players. Each group of bars shows the average ϵ for equilibria of reductions of the given game at increasing fidelity. The number of players in the full game is shown in parentheses, with the number in reduced games under each bar. The bars extend upward to indicate a 95% confidence upper bound on ϵ. To the left of each group is shown the ϵ (with 95% confidence interval) of the social optimum of the full game.

Strategies include all policies for bidding on flights, hotels, and entertainment over time, as a function of prior observations. The agents interact in the markets for each kind of good, as competing buyers or potential trading partners. Based on published accounts, TAC participants design agents given specified game rules, and then test these designs in the actual tournaments as well as offline experiments. Testing is crucial, given the lack of any compact analytical model of the domain. In experiments, agent designers explore variations on their agent program, for example by tuning parameters or toggling specific agent features.

That strategic choices interact, and implications for design and evaluation, have been frequently noted in the TAC literature. We omit a detailed catalog due to space limitations, mentioning only the extensive experimental effort on TAC travel-shopping reported by Vetsikas and Selman (2003). In the process of designing Whitebear for TAC-02, they first identified candidate policies for separate elements of the agents' overall strategy. They then defined extreme (boundary) and intermediate values for these partial strategies, and performed experiments according to a systematic and deliberately considered methodology. Specifically, for each run, they fix a particular number of agents playing intermediate strategies, varying the mixture of boundary cases across the possible range. In all, the Whitebear experiments comprised 4500 profiles, with varying *even* numbers of candidate strategies (i.e., profiles of the 4-player game). This systematic exploration was apparently helpful, as Whitebear was the top scorer in the 2002 tournament. This agent's predecessor version placed third in TAC-01, following a less comprehensive and structured experimentation process. Its successor placed third again in 2003, and regained its first-place standing in 2004.

TAC Experiments

To apply reduced-game analysis to the TAC domain, we identified a restricted set of strategies, defined by setting parameters for Walverine (Cheng *et al.* 2005). We considered a total of 35 distinct strategies, covering variant poli-

cies for bidding on flights, hotels, and entertainment. A description of this parametrization and its use in the design of our 2005 TAC entry are provided elsewhere (Wellman *et al.* 2005). We collected data for a large number of games: over 37,000 as of this writing, representing over ten months of (almost continuous) simulation.[4] Each game instance provides a sample payoff vector for a profile over our restricted strategy set.

Table 1 shows how our dataset is apportioned among the 1-, 2-, and 4-player reduced games. We are able to exhaustively cover the 1-player game, of course. We could also have exhausted the 2-player profiles, but chose to skip some of the less promising ones (around one-quarter) in favor of devoting more samples elsewhere. The available number of samples could not cover the 4-player games, but as we see below, even 2.4% is sufficient to draw conclusions about the possible equilibria of the game. Spread over the 8-player game, however, 37,000 instances would be insufficient to explore much, and so we refrain from any sampling of the unreduced game.

p	Profiles			Samples/Profile	
	total	evaluated	%	min	mean
4	73,815	1775	2.4	10	20.8
2	630	467	74.1	15	31.1
1	35	35	100.0	20	91.5

Table 1: Profiles evaluated, reduced TAC games (TAC$\downarrow_p$).

In the spirit of hierarchical exploration, we sample more instances per profile as the game is further reduced, obtaining more reliable statistical estimates of the coarse background relative to its refinement. On introducing a new profile we generate a minimum required number of samples, and subsequently devote further samples to particular profiles based on their potential for influencing our game-theoretic analysis. The sampling policy employed was semi-manual and somewhat *ad hoc*, driven in an informal way by analyses of the sort described below on intermediate versions of the dataset. Developing a fully automated and principled sampling policy is the subject of future research.

1-Player Game

The 1-player game (TAC$\downarrow_1$) would typically not merit the term "game", as it assumes each strategy plays only among copies of itself. Thus, its analysis considers no strategic interactions. To "solve" the game, we simply evaluate which has the greatest expected payoff. For our experiment, we obtained 20–267 samples of each of the 35 1-player profiles, one for each strategy.

Figure 5 displays the average payoffs for each 1-player profile, sorted from best-to-worst, left-to-right. We tended to take more samples of the more promising profiles, but

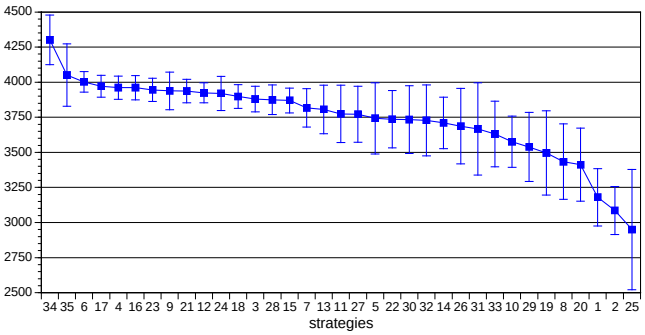


Figure 5: Average payoffs for strategy profiles in TAC$\downarrow_1$. Error bars delimit 95% confidence intervals.

cannot statistically distinguish every profile in the ranking. Nevertheless our top strategy, number 34,[5] performs dramatically—250 points—better than the next best, number 35. Pairwise mean-difference tests rank 35 above all others at significance levels of $p < 0.05$.

In the absence of further data, we might propose strategy 35, the unique pure-strategy Nash equilibrium (PSNE) of the 1-player game. In fact, however, this strategy was designed expressly to do well against itself, and may be vulnerable in environments with other agents. By exploring a less extreme reduction we can start to consider some of the strategic interactions.

2-Player Game

The two-player game, TAC$\downarrow_2$, comprises 630 distinct profiles: $35 \cdot 34/2 = 595$ where two different strategies are played by four agents each, plus the 35 profiles from TAC$\downarrow_1$ where all agents play the same. We can identify PSNE simply by examining each strategy pair (s, s'), and verifying whether each is a best response to the other. In doing so, we must account for the fact that our sample data may not include evaluations for all possible profiles.

Definition 4 *Profiles can be classified into four disjoint categories, defined below for the 2-player pure-strategy case. (The generalization to N-player is straightforward.)*

1. *If (s, s') has not been empirically evaluated, then $\hat{u}(s, s')$ is undefined, and we say (s, s') is* unevaluated.
2. *Otherwise, and for some t, $\hat{u}(t, s') > \hat{u}(s, s')$ or $\hat{u}(t, s) > \hat{u}(s', s)$. In this case, we say (s, s') is* refuted.
3. *Otherwise, and for some t, (t, s') is unevaluated or (s, t) is unevaluated. In this case, we say (s, s') is a* candidate.
4. *Otherwise, in which case we say (s, s') is* confirmed.

Based on our TAC$\downarrow_2$ simulations, we have confirmed five PSNE: (3,23), (4,9), (5,16), (6,17), and (7,24). We have refuted 462 profiles, and the remaining 163 are unevaluated.

The definitions above say nothing about the statistical strength of our confirmation or refutation of equilibria. For any particular comparison, one can perform a statistical

[4] Our simulation testbed comprises two dedicated workstations to run the agents, another RAM-laden four-CPU machine to run the agents' optimization processes, a share of a fourth machine to run the TAC game server, and background processes on other machines to control the experiment generation and data gathering.

[5] Since our present purpose is to demonstrate techniques for exploration and strategic reasoning rather than to address the substance of trading strategies, we identify them by index only.

analysis to evaluate the weight of evidence for or against stability of a given profile. For instance, we could construct diagrams of the form of Figure 5, but representing the payoff in response to a particular strategy, rather than in self-play. Such a plot of responses to strategy 17 would indicate, for example, that 18 is quite nearly as good as 6, and so the confirmation of (6,17) as a PSNE is statistically weak.

We can also measure the *degree* of refutation in terms of the ϵ measure defined by (1). Since the payoff function is only partially evaluated, for any profile we have a *lower bound* on ϵ based on the deviation profiles thus far evaluated. We can generalize the classifications above (refuted, candidate, confirmed) in the obvious way to hold with respect to any given ϵ level. For example, profile (17,18) is confirmed at $\epsilon = 0.08$, but all other non-PSNE profiles are refuted at $\epsilon > 13$. Figure 6 presents the distribution of ϵ levels at which the 467 evaluated 2-player profiles have been refuted. For example, over half have been refuted at $\epsilon > 265$, and all but 10 at $\epsilon > 30$. These 10 pure profiles remain candidates (9 of them confirmed) at $\epsilon = 30$.

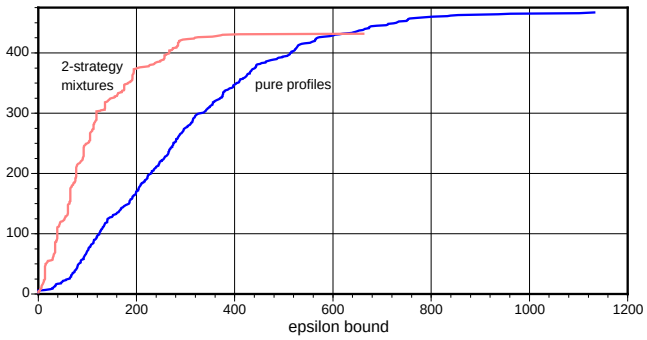


Figure 6: Cumulative distribution of ϵ bounds in TAC$\downarrow_2$.

We can also evaluate symmetric profiles by considering mixtures of strategies. Although we do not have the full payoff function, we can derive ϵ bounds on mixed profiles, as long as we have evaluated pure profiles corresponding to all combinations of strategies supported in the mixture. For example, we can derive such bounds for all 432 *pairs* of strategies for which we have evaluated 2-player profiles. The distribution of bounds for these pairs are also plotted in Figure 6. Note that the ϵ bound for a strategy pair is based on the *best* mixture possible of that pair, and so the refutation levels tend to be smaller than for pure strategies. Indeed, three pairs—(4,9), (5,16), (6,17)—participate in confirmed equilibria, another—(34,35)—is a candidate, and a total of 17 pairs remain candidates at $\epsilon = 10$, with 12 confirmed at that level.

We apply the term k-*clique* to a set of k strategies such that all profiles involving these strategies are evaluated. A clique defines a subgame of the original game, which can be evaluated by standard methods. We applied iterative elimination of dominated strategies to all the maximal cliques of the 2-player game, ranging in size up to $k = 23$. This indeed pruned many strategies and induced new subsumption relations among the cliques, leaving us with only one maximal clique, of size 16. We applied the Lemke-Howson algorithm to this subgame, which identified 29 candidate symmetric equilibria (not refuted by strategies outside the cliques), with distinct supports ranging in size from two to nine. Nineteen of these mixtures are confirmed (including the three pairs mentioned above).

Because any equilibrium of the full game must also be an equilibrium in any subgame encompassing its support, this exercise also allows us to prune broad regions of profile space from consideration.[6] For instance, the subgame results effectively refute 3056 strategy triples (out of 6545 total, or 47%) as comprising support for symmetric equilibria. By similar reasoning, we refute 14789 strategy quadruples (28%). Given the importance of small supports in recent approaches to deriving equilibria (Porter, Nudelman, & Shoham 2004), or approximate equilibria (Lipton, Markakis, & Mehta 2003), focusing the search in these regions can be quite helpful.

Finally, we can account for statistical variation in the estimated payoffs by employing sensitivity analysis in our ϵ calculations. Specifically, we interpret each payoff value in the estimated game as normally distributed with mean and variance given by the sample. We then apply Monte Carlo methods to generate a distribution of ϵ values for a given profile, one corresponding to each draw of a payoff function from the specified distributions. Naturally, even our confirmed equilibria are refuted with substantial probability, and thus have positive ϵ in expectation. The most robustly stable profile we have identified thus far is a mixture of (3,16,23), with a mean ϵ value of 64.

4-Player Game

Our analysis of the 4-player game, TAC$\downarrow_4$, parallels that of the 2-player game, though of course based on a sparser coverage of the profile space. There are 73,815 distinct TAC$\downarrow_4$ profiles, out of which we have evaluated 1775. Of these, 154 are TAC$\downarrow_2$ profiles with no evaluated neighbors in TAC$\downarrow_4$ (i.e., no deviations tested). Although these are technically PSNE candidates, we distinguish them from the one PSNE candidate that has actually survived some challenge. The remaining 1620 evaluated profiles are refuted, at various levels. The distribution of ϵ bounds is plotted in Figure 7.

Figure 7 also shows, inset, the distribution of epsilon bounds over the 128 strategy pairs for which we have evaluated all combinations in TAC$\downarrow_4$ (i.e., the 2-cliques). Among these are 14 candidate equilibria at $\epsilon = 10$, two of them nearly confirmed at $\epsilon = 0.1$. The TAC$\downarrow_4$ cliques are relatively small: three 5-cliques, 15 4-cliques, and 51 3-cliques. Eliminating dominated strategies prunes little in this case, and we have been unsuccesful in getting GAMBIT to solve any k-clique games in the 4-player game for $k > 2$. However, applying replicator dynamics produces sample symmetric subgame equilibria, including 11 mixture triples that constitute candidates with respect to the full game.

[6] Pruning is strictly justified only under the assumption that we have identified *all* symmetric equilibria of the clique subgames. The Lemke-Howson algorithm does not guarantee this, but in every case for which we were able to check using more exhaustive methods available in GAMBIT, in fact all such equilibria were found.

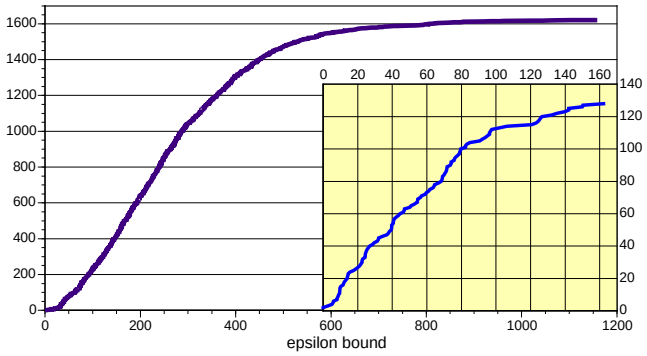


Figure 7: Cumulative distribution of ϵ bounds in TAC$\downarrow_4$. Main graph: pure profiles. Inset: 2-strategy mixtures.

Finally, given data in both the 2-player and 4-player games, we can perform some comparisons along the lines of our GAMUT experiments described above. The results, shown in Figure 8, are not as clear as those from the known-game experiments, in part because there is no "gold standard", as the 4-player game is quite incompletely evaluated.

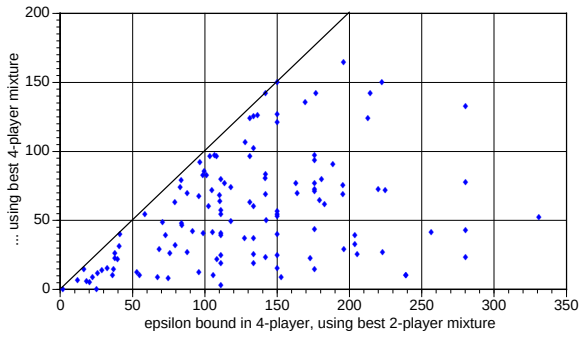


Figure 8: ϵ bounds in the 4-player game achieved by playing the best mixture from the 2-player game, versus playing that best in the 4-player. All points must be southeast of the diagonal by definition.

Discussion

Given all this simulation and analysis, can we now identify the "best" strategy to play in TAC? Surely not, though we do have strong evidence for expecting that a sizable fraction (about two-thirds) of the original 35 strategies will turn out to be unstable within this set. However, we still lack a definitive characterization of equilibrium profiles for the game, and moreover even possessing one would not necessarily give us the answer. Strategic stability is just one form of evidence bearing on what we should expect other agents to play. In lieu of a categorical recommendation, what we have is an improved understanding of the strategic landscape. Perhaps more importantly, given a new strategy proposal, our analysis tells us in which strategic contexts (i.e., the relatively stable profiles) it needs to be evaluated. We have adopted this heuristic in exploring the strategy space for Walverine, and will soon learn how well this approach prepared us for the 2005 tournament (Wellman *et al.* 2005).

More generally, we conclude that hierarchical analysis of reduced games can be an effective tool for scaling up empirical game-theoretic methods to symmetric environments with many agents. Choice of reduction level trades fidelity for tractability in a controlled way. Further theoretical analysis and application experience should lead to a fuller understanding of the tradeoff, and identify additional ways to apply the reduction approach presented here.

Acknowledgments

We thank Eugene Nudelman and Ted Turocy for assistance with GAMUT and GAMBIT, respectively. This work was supported in part by the National Science Foundation under grants IIS-0205435 and IIS-0414710, and by the DARPA REAL program.

References

Cheng, S.-F.; Leung, E.; Lochner, K. M.; O'Malley, K.; Reeves, D. M.; and Wellman, M. P. 2005. Walverine: A Walrasian trading agent. *Decision Support Systems* 39:169–184.

Krishna, V. 2002. *Auction Theory*. Academic Press.

Leyton-Brown, K., and Tennenholtz, M. 2003. Local-effect games. In *Eighteenth International Joint Conference on Artificial Intelligence*, 772–780.

Lipton, R. J.; Markakis, E.; and Mehta, A. 2003. Playing large games using simple strategies. In *Fourth ACM Conference on Electronic Commerce*, 36–41.

Mas-Colell, A.; Whinston, M. D.; and Green, J. R. 1995. *Microeconomic Theory*. Oxford University Press.

McKelvey, R. D.; McLennan, A.; and Turocy, T. 1992. Gambit game theory analysis software and tools. `http://econweb.tamu.edu/gambit`.

Nash, J. 1951. Non-cooperative games. *Annals of Mathematics* 54:286–295.

Nudelman, E.; Wortman, J.; Shoham, Y.; and Leyton-Brown, K. 2004. Run the GAMUT: A comprehensive approach to evaluating game-theoretic algorithms. In *Third International Joint Conference on Autonomous Agents and Multi-Agent Systems*, 880–887.

Porter, R.; Nudelman, E.; and Shoham, Y. 2004. Simple search methods for finding a Nash equilibrium. In *Nineteenth National Conference on Artificial Intelligence*, 664–669.

Vetsikas, I. A., and Selman, B. 2003. A principled study of the design tradeoffs for autonomous trading agents. In *Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*, 473–480.

Wellman, M. P.; Greenwald, A.; Stone, P.; and Wurman, P. R. 2003. The 2001 trading agent competition. *Electronic Markets* 13:4–12.

Wellman, M. P.; Reeves, D. M.; Lochner, K. M.; and Suri, R. 2005. Searching for Walverine 2005. In *IJCAI-05 Workshop on Trading Agent Design and Analysis*.