

Best-First Search for Approximate Equilibria in Empirical Games

Patrick R. Jordan and Michael P. Wellman

University of Michigan
Computer Science and Engineering
2260 Hayward
Ann Arbor, MI 48109-2121
{prjordan, wellman}@umich.edu

Abstract

When exploring a game over a large strategy space, it may not be feasible or cost-effective to evaluate the payoff of every relevant strategy profile. For example, evaluating each payoff of an empirically defined game may require Monte Carlo simulation or other costly computation. Analyzing such games poses a *search problem*, with the goal of identifying and confirming pure-strategy equilibrium profiles by evaluating payoffs of candidates and potential deviations from those candidates. Sureka & Wurman (2005) studied this problem and proposed a search method based on best-response dynamics. We introduce a family of *best-first* algorithms, which prioritize unconfirmed profiles by their known bound away from equilibrium and select the profile with the current minimum. We compare algorithms by measuring the fraction of profile space explored to confirm equilibria, as well as the search effort required to confirm approximate equilibria, on several game classes. Our best-first approach compares similarly to the existing best-response algorithms when searching for exact equilibria, and favorably when searching for approximate equilibria.

Introduction

In attempting to understand agent interactions in multiagent systems, researchers often appeal to game-theoretic solution concepts in characterizing the strategic stability of hypothetical outcomes. Unfortunately, the strategy space of the game or interaction being modeled is often so complex to render infeasible exact game-theoretic modeling and analysis. One common compromise is to consider stylized versions of the game that are amenable to computational analysis, at the expense of fidelity. One alternative pursued by experimental AI researchers in recent years is to estimate games through simulation and sampling, an approach that has been termed *empirical game-theoretic analysis* (Reeves, 2005; Wellman, 2006). This approach has held particular appeal for researchers interested in trading agent games, including ourselves as well as the authors of much of the related work cited herein.

In empirical game modeling, the outcome of a joint strategy, or *profile*, is estimated by repeatedly sampling the game. These samples can be generated by a game simulator or other model describing the game. Such an approach is quite general, but incurs an estimation cost in proportion to the size of the profile space, which is exponential in the

number of players and the number of strategies available per player. For many games of interest, the strategy set is extremely large or infinite. Thus, in practice we cannot explore the space exhaustively, but instead focus on profiles that are most promising as solutions or otherwise pivotal in game-theoretic analysis. The choice of profiles to explore is essentially a process of *heuristic search*, and effective heuristics are essential to computationally feasible strategic analysis in empirical games.

Figure 1 gives an overview of empirical game-theoretic analysis. The payoff function mapping profiles to payoffs, which is found in a typical game definition, has been replaced by a simulator and a set of sample observations. Using this framework, agent designers can construct and analyze heuristic strategies from a large strategy space (Kiekintveld, Wellman, & Singh, 2006), and mechanism designers can optimize an objective function whose computation was previously intractable (Vorobeychik, Kiekintveld, & Wellman, 2006). For both of these design problems, it is useful to determine the set of exact or approximate Nash equilibria, which thus constitutes a core search problem for empirical strategy design and empirical mechanism design.

Previous research has explored directed sampling of profiles (Walsh, Parkes, & Das, 2003; Reeves *et al.*, 2005) by using value of information estimates and interleaving sampling and equilibrium calculations, respectively. Both techniques require at least a small number of samples to be generated for each of the profiles in the full joint strategy space. Since it may be possible to establish that a particular profile is an equilibrium or near-equilibrium without considering all profiles, a search approach can potentially relax this requirement. This was part of the motivation for Sureka & Wurman (2005) to develop an algorithm, based on tabu best-response search, to search for pure Nash equilibria within the profile space.

While search strategies such as tabu best-response allow for partial sampling of the profile space, they do not exploit information contained in the interactions of strategies of partially specified profiles. We have encountered many games where particular combinations of strategies predictably do not interact well regardless of the remaining players' strategies. This type of substructure can be exploited to prune the search space. In general, predictive models for the best deviation of a particular partial or full profile can be constructed

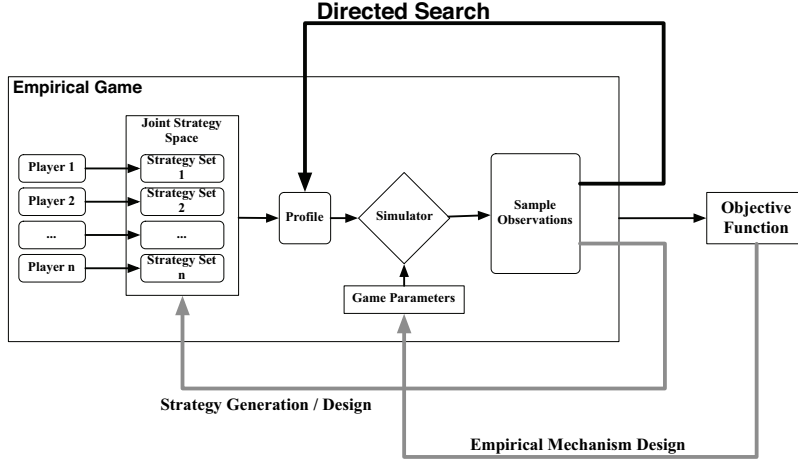


Figure 1: Overview of empirical game-theoretic analysis where directed search is used to reduce the number of observations sampled in confirming profiles with low ϵ . These profiles are extensively used in strategy and mechanism design.

to attempt to improve directed search and allow more interesting regions of the profile space to be explored.

Background

This section describes the formal notation and experimental measures of success we use to compare algorithms. We follow standard game-theoretic notation, but consider only the pure strategy space of players within a game. Our first performance measure is a variant of the measure described by Sureka & Wurman (2005), and the second extends this to consider approximate equilibria. Approximate equilibria may be of interest in general, and are especially salient when searching among pure profiles for games that may not exhibit pure-strategy Nash equilibria.

Formal Notation

Definition 1 (Normal Form Game). $\Gamma = \langle I, \{S_i\}, \{u_i(s)\} \rangle$ is a normal form game where I is the set of players, S_i is the set of strategies available to player i , and $u_i : \times_{j=1}^{|I|} S_j \rightarrow \mathbb{R}$ is the utility function for player i mapping the joint strategy s to the real-valued payoff received by player i when s is played.

We refer to the joint strategy set $\times_{j=1}^{|I|} S_j$ for a particular game Γ as $S(\Gamma)$ and we drop the Γ parameterization when the context is clear. The term *profile* is used interchangeably with *joint strategy*.

Definition 2 (Unilateral Deviation Set). For some joint strategy $s \in S(\Gamma)$, the unilateral deviation set $\mathcal{D}_i(s)$ is

$$\mathcal{D}_i(s) = \{(\hat{s}_i, s_{-i}) : \hat{s}_i \in S_i - \{s_i\}\} \quad (1)$$

and

$$\mathcal{D}(s) = \bigcup_{i \in I} \mathcal{D}_i(s). \quad (2)$$

Definition 3 (ϵ of a joint strategy). For some joint strategy $s \in S(\Gamma)$, the maximum gain from deviation, $\epsilon(s)$, is

$$\epsilon(s) = \max_{i \in I, \hat{s} \in \mathcal{D}_i(s) \cup \{s\}} u_i(\hat{s}) - u_i(s). \quad (3)$$

Definition 4 (Pure Strategy Nash Equilibrium). For some joint strategy $s \in S(\Gamma)$, s is a pure strategy Nash equilibrium (PSNE) iff $\epsilon(s) = 0$.

Repeated Sampling Approach

In games with stochastic payoffs multiple samples may be required to accurately estimate the expected value of a profile. Generally, solutions to this problem have been of the form

- (a) Gather an initial sample for each possible profile
- (b) Update the payoff estimates
- (c) Repeatedly sample interesting profiles
- (d) Go to (b)

Walsh, Parkes, & Das (2003) approach the sampling problem by using the value of information framework developed by Russell & Wefald (1991). The idea of this approach is to select the sample s that is expected to provide the greatest reduction in estimated error in the equilibrium choice. In the original setting, sample s is selected in an attempt to refute an equilibrium. Since this computation is very costly, Walsh, Parkes, & Das (2003) propose an alternate computation which selects profiles based on the *confirmational value of information*.

A second approach is due to Reeves *et al.* (2005). The algorithm uses replicator dynamics to select samples from an initially uniform distribution. This form of biased sampling assigns higher probability of sampling to profiles with more density in the replicator dynamics mixture. In essence, the hope is that more important samples will be selected more often so that we can obtain accurate equilibrium mixtures without repeatedly sampling less important regions of the profile space.

As previously mentioned, these approaches require that the entire space has been partially sampled. In many games of interest this is prohibitively expensive. For games amenable to search with best-response dynamics, we show

empirically that pure-strategy Nash equilibria are typically confirmed with less than 1% of the space searched. If an approach is taken such that each profile searched is sampled repeatedly, say 30 times, for statistical significance tests, this results in a sample count equal to approximately 30% of the search space. This is at worst a 70% savings over just the initial sampling step in the procedure sketched above.

In the next subsection we discuss an alternative method proposed by Sureka & Wurman (2005) which does not require the entire profile space to be searched. Since only pure profiles are considered, we lose the guarantee that an exact Nash equilibrium will be found.

Tabu Best-Response Search

The procedure introduced by Sureka & Wurman (2005) samples all deviations of the active profile for a given player, and selects the best response (i.e., the profile maximizing gain from deviation) as the next active profile. The algorithm then selects a new player and the process iterates. To avoid best-response cycles, they maintain a *tabu list* recording profiles that have already been visited. Two variants of the algorithm were discussed: one with *explicit memory* and another with *attribute based memory*. The explicit memory version keeps a global tabu list holding all previously visited joint strategies, whereas attribute-based memory keeps a separate tabu list for each player. The process terminates once the algorithm selects a PSNE as its active profile.

The original experiments by Sureka & Wurman (2005) focus on the time required to find a PSNE. The equilibria of the games were pre-computed, therefore finding a PSNE amounts to checking the active profile against a list of known equilibria. We would like to consider games in which the equilibria are not externally verified and, thus, the search procedure must itself confirm their existence. Therefore, in order to compare how well the algorithms confirm equilibria, we slightly modify tabu best-response search. Now, instead of immediately placing the active profile on the tabu list T and branching to the best response, we branch to the best response only if that profile is improving. Of course, no player would choose to deviate in a Nash equilibrium, and by keeping the profile active we allow it to be confirmed by iteration through the players. With this modification, however, it becomes possible that we visit a profile for which all neighbors are in the tabu list (explicit memory version). In this case we allow the player to deviate to the best response if that best response gives a higher payoff than the current profile. Pseudo-code for the tabu best-response algorithm used in our experiments is presented below.

TABU-BEST-RESPONSE-SEARCH(Γ)

```

Select initial profile at random
while Termination criteria not satisfied
  do
     $i \leftarrow$  next player
    if  $D_i(s) \subseteq T$ 
      then  $s \leftarrow$  player  $i$ 's best response to  $s$ 
    else if  $s$  has an improving deviation in  $D_i(s) \setminus T$ 
      then Push  $s$  onto  $T$ 
       $s \leftarrow$  player  $i$ 's best response to  $s$  not in  $T$ 

```

Performance Measure

We wish to compare best-first search (BFS) against tabu best-response (TABU) in a test of search efficiency. Our preference is for a search algorithm which confirms low ϵ profiles sooner. When testing TABU, Sureka & Wurman (2005) measured performance by the percentage of strategy space examined to *find* a sample Nash equilibrium. To directly compare TABU to BFS we measure performance in terms of the percentage of the strategy space examined to *confirm* a sample Nash equilibrium. To confirm an equilibrium, all of the profile's deviations must be sampled.

The unfortunate consequence of using the first measure is that we are restricting our comparison to games with at least one PSNE. In general we are concerned with games that may not have a PSNE (at least with respect to the strategy space considered), and even when such exist we may also be interested in approximate equilibria. Therefore we introduce a second performance measure which characterizes the lowest confirmed ϵ as a function of the percentage of profile space searched.

Best-First Search

The idea of BFS is to expand a search tree by exploring the fringe node that is best according to some priority measure. In our setting, the objective is to find a profile minimizing potential gain from deviation, ϵ . Therefore we adopt as our priority measure a lower bound, $\hat{\epsilon}(s)$, on the possible gain to deviation from profile s . The pseudo-code below describes the BFS procedure as applied to game Γ .

BEST-FIRST-SEARCH(Γ)

```

Select initial profile at random
while Queue is not empty
  do
    Select lowest  $\hat{\epsilon}(s)$  profile  $s$  from queue
    if  $s$  is confirmed
      then Remove it from queue and assign  $\epsilon(s) = \hat{\epsilon}(s)$ 
    else  $\bar{s} \leftarrow$  SELECT-DEVIATION( $s$ )
      Insert  $\bar{s}$  into queue if previously unsampled
      Update  $\hat{\epsilon}(\bar{s})$  for  $\hat{s} \in \{\bar{s}\} \cup \mathcal{D}(\bar{s})$  in the queue

```

The lower bound $\hat{\epsilon}(s)$ is simply the maximal gain found among the deviations from profile s sampled thus far. A profile s is *confirmed* if all the deviations from s have been sampled. SELECT-DEVIATION(s) returns an unsampled deviation from s . Selection is done by predicting which unsampled deviation from s is likely to give the largest gain from deviation. While the efficacy of the deviation selection heuristic will depend on the class of games being tested, we have empirically found one heuristic which works well in a variety of cases (Jordan, Kiekintveld, & Wellman, 2007). Consequently, the prediction model used in this paper will calculate the percentage of historically improving deviations and pick the deviation which maximizes the improving deviation probability.

As with any heuristic search, BFS can be hindered by unfavorable terrain. Figure 2 illustrates how this can occur

within the joint strategy space of a game. In this case, the four profiles with high ϵ values (relative to the neighboring profiles) are blocking the profile subspace that contains the optimal low ϵ profile. In BFS the low ϵ profile will not be confirmed until exploration of one of the four blocking profiles, unless the initial profile of the search happens to be one of the five labeled profiles.




Figure 2: Blocking an optimal subspace. Profiles with high ϵ are labeled **H** and the optimal profiles with low ϵ is labeled **L**. Arrows represent position gain unilateral deviations.

When there is a known upper bound on the minimum ϵ we can employ random restart to mitigate this kind of local blocking problem. The standard BFS algorithm chooses the next profile from among the unsampled deviations of the profile with the lowest current ϵ . Using the known bound on ϵ , we uniformly select a random profile over all unsampled profiles if the $\hat{\epsilon}$ on the best profile is above this bound. We denote this algorithm BFS-RR and use the bound of zero for experimentation.

Random Games with Independent Payoffs

We tested various classes of games using GAMUT (Nudelman *et al.*, 2005). Of those classes of games, one which received substantial attention by Sureka & Wurman (2005) is the class with payoffs that are independently identically distributed. Sureka & Wurman (2005) chose this class of game to benchmark TABU, due to the lack of inherent structure. We follow suit by presenting an algorithm specifically designed to exploit this class, which we use as a benchmark for BFS and TABU. We label this algorithm frequency first search (FFS) for reasons that will shortly become apparent.

Notice that by the assumption of independent payoffs,

$$\mathbb{P}(\epsilon(s) = 0) = \prod_{i \in I, \hat{s} \in D_i(s)} \mathbb{P}(u_i(\hat{s}) \leq u_i(s)).$$

Therefore, for a given joint strategy s that has been sampled, we can calculate the probability that s is a PSNE. In this case $u_i(\hat{s})$ is either a random variable and the probability that the deviation gain is not positive is computed by evaluating the cumulative distribution function of the payoffs, or the deviation profile has been sampled and we can assign the probability a 0 or a 1. Similarly, for an unsampled profile we

can determine the probability by evaluating the cumulative distribution function at the sampled deviation payoff. In the case that the density function of the payoffs is not known we can approximate the calculation by using the sample distribution. FFS will similarly perform best-first search, but rather than ordering by $\hat{\epsilon}$, FFS will instead order by probability of being a PSNE.

Experimental Results

Our experiments use games of various classes generated by GAMUT (Nudelman *et al.*, 2005) as well as a set of empirical games in a trading agent domain. When applicable, games similar in size to those given by Sureka & Wurman (2005) are tested. Initially we experiment with the game classes use in this prior study to establish a baseline for algorithm comparison. We then proceed to evaluate various game classes whose structure is known to be exploited by the best-response dynamics so that we may test BFS in an environment that should be favorable to TABU. Finally, we conclude with our analysis of the algorithms by analyzing their performance on a set of TAC/SCM games.

Uniform Random Games

Our first class of games have payoffs that are uniformly distributed in the range $[-100,100]$ denoted $URG(|I|, |S_i|)$ where $|I|$ is the number of players and $|S_i|$ is the strategy set size of each player. We compare BFS and TABU on two sizes of games: the smaller $URG(5,5)$ and the larger $URG(5,10)$. To construct the data sets for comparing the algorithms we systematically generated the games and checked whether these games contained a PSNE. Games that fell within this criterion were tested using the first measure of efficiency (the % of profile space searched to confirm an equilibrium), as well as, the second. For $URG(5,5)$ the results of this comparison are shown in Table 1.

	BFS	FFS	Tabu
Mean (%)	53.42	34.36	52.25
Median (%)	52.12	33.40	49.28
BFS	–	2.2e-16	0.18
FFS		–	2.2e-16
Tabu			–

Table 1: $URG(5,5)$ with at least one PSNE. Statistics for each method and p-value comparisons.

Our analysis of $URG(5,5)$ included 12 games which contained at least one PSNE. Seeds 0, 1, 3, 4, 8, and 15 contained one PSNE; seeds 5, 13, 16, 17, and 18 contained two; and seed 20 contained three. FFS was a statistically significant improvement over both BFS and TABU whose difference in the first measure was not significant. In each game 100 profiles were selected at random as starting points for all algorithms. The performance of BFS and TABU varied drastically according to the individual game. For instance, in seeds 4, 8, and 15 TABU rarely succeeded in confirming the solution.¹ Similarly in seed 3, BFS on average requires

¹Since TABU is not guaranteed to confirm an existing solution a

nearly all the search space to be evaluated. It should be noted that although the global optimal was not confirmed until near the last iteration, many near optimal profiles were confirmed much earlier.

Figure 3 shows the minimum confirmed ϵ as a function of the space explored, which is our second performance measure. In many practical settings it may be that near equilibrium profiles are just as useful as PSNE. Therefore in those cases we consider the second measure more appropriate. Notice that BFS confirms low ϵ profiles much earlier than TABU, which is the desired result.




Figure 3: Mean low ϵ for URG(5,5).

	BFS	BFS RR	FFS	Tabu
Mean (%)	37.10	35.68	23.88	41.79
Median (%)	31.41	28.88	20.31	34.75
BFS	–	0.19	2.2e-16	3.5e-05
BFS RR		–	2.2e-16	6.1e-08
FFS			–	2.2e-16
Tabu				–

Table 2: URG(5,10) with at least one PSNE. Statistics for each method and p-value comparisons.

Our analysis of URG(5,10) also included 12 games which contained at least one PSNE. All algorithm comparisons are significant with the exception of the BFS and BFS-RR algorithms. In these larger games the average performance of BFS and TABU improves from approximately 50% of the space searched to the mid-30% range. Here again, as expected, FFS produces the best performance with only 24% of the space searched on average. Figure 4 shows the minimum confirmed ϵ as a function of the space explored for the large game. BFS shows a large improvement over TABU in the large game as well.

Covariance Games

In covariance games payoffs are distributed normally with zero mean, unit variance, and covariance r between player's

timeout was placed on the number of iterations equal to the size of the strategy space. If TABU exceeded the timeout it was credited for finding the solution in the greatest possible number of steps required by BFS.




Figure 4: Mean low ϵ for URG(5,10).

payoffs in the same profile. The parameter r used by GAMUT to create the range is randomly assigned according to GAMUT's parameter randomization routine. We generate 20 instances of this game type with 5 players and 10 actions each and report the results of analyzing the algorithms according to the first efficiency measure.

	BFS	Tabu
Mean (%)	0.22	0.36
Median (%)	0.17	0.15
p value	0.10	

Table 3: Covariance Game (5,10). Statistics for each method and p-value comparisons.

Table 3 gives the results of the comparisons on covariance games. The 0.36% mean percentage of space searched by TABU is much lower than the approximate 1.1% given by Sureka & Wurman (2005).² This can probably be explained by the fact that Sureka & Wurman (2005) ran experiments initialized with each profile in the joint profile space. We randomly select 100 profiles to initialize the search for each algorithm and while some PSNE may have large basins of attraction which result in quick convergence, there may exist a small region of the space which takes much longer to search resulting in a low probability, large value event. As with the uniform random games, neither algorithm shows a significant improvement over the other.

Congestion Games

Here we repeat the description of congestion games generated by GAMUT given in the user documentation:

In the congestion game, each player chooses a subset from the set of all facilities. Each player then receives a payoff which is the sum of payoff functions for each facility in the chosen subset. Each payoff function

²Sureka & Wurman (2005) used a 5 player 11 action version of the game.

depends only on the number of other players who have chosen the facility.

The class of congestion games have been shown to be equivalent to the class of potential games (Rosenthal, 1973; Monderer & Shapley, 1996), which are games that have payoffs given by a potential function. Congestion games have a very special property proved by Rosenthal (1973): *congestion games always have a PSNE*.

We compare BFS and TABU on four-player four-facility congestion games. The results of the congestion game comparison are shown in Table 4. Twenty games were generated for experimentation using GAMUT. In all cases TABU has better performance than BFS. This is reflected in the results table for this class of games, but it should be noted that the difference in space explored, 0.05%, is negligible.

	BFS	Tabu
Mean (%)	0.15	0.10
Median (%)	0.15	0.10
<i>p</i> value	< 2.2e-16	

Table 4: Congestion game (4,4). Statistics for each method and *p*-value comparisons.

Local Effect Games

Local effect games (LEGs) (Leyton-Brown & Tennenholtz, 2003) are symmetric games in which pairs of actions locally affect each other if the utility of agent taking one of the actions depends on some function of the number of agents taking the other action. Leyton-Brown & Tennenholtz (2003) find experimentally that myopic best-response dynamics converge quickly to PSNE in sample LEGs. Given this prior evidence we expect TABU to measure well on these games and this expectation is confirmed in our experiments shown in Table 5. As with the other experiments, we generate 20 instances of this game type with 5 players and 10 actions each, then report the results of analyzing the algorithms according to the first efficiency measure.

	BFS	Tabu
Mean (%)	1.52	0.09
Median (%)	0.12	0.08
<i>p</i> value	7e-16	

Table 5: LEG (5,10). Statistics for each method and *p*-value comparisons.

Table 5 gives the results of the analysis. TABU has a significant edge over BFS in both the mean and median of the percentage of space searched. The relatively large mean of BFS is due to seed 14 of the experiments. In this game BFS did not find the PSNE until approximately 40% of the space was searched.

TAC/SCM_{↓3} Games

We conclude our analysis by comparing the algorithms on an empirical version of the Trading Agent Competition

Supply Chain Management (TAC/SCM) game. TAC/SCM (Arunachalam & Sadeh, 2005; Eriksson, Finne, & Janson, 2006) is a six-player game in which agents representing PC (personal computer) manufacturers compete to maximize their profits over a simulated year. There are 220 scenario days, and agents have approximately 14 seconds to make decisions each day. Agents participate simultaneously in markets for supplies (components) and finished PCs. There are 16 different types of PCs (divided into three market segments), defined by the compatible combinations of 10 different component types. Components fall into one of four categories: CPU, motherboard, memory, and hard disk. There are four types of CPUs and two types of all other components; one component from each category is required to produce a PC.

The strategy space for our empirical version of game (Jordan, Kiekintveld, & Wellman, 2007) comprises a subset of the agents participating in the 2005 and 2006 TAC/SCM competitions. Binary versions of these agents, listed in Table 6, were provided to the TAC agent repository by their designers. We employed hierarchical game reduction (Wellman *et al.*, 2005) to approximate the six-player game by a three-player version, called SCM_{↓3}. Each player in SCM_{↓3} is represented by two copies of its chosen strategy in the original six-player game. We estimated the payoffs for two empirical SCM_{↓3} games, corresponding to the 2005 and 2006 agents, through repeated sampling of each distinct profile of the respective strategy sets.

Agent	Tournament Scores Finals
Deep Maize 05	-0.22
Go Blue Oval 05	n/a
Mertacor 05	0.55
MinneTAC 05	-0.31
PhantAgent 05	n/a
TacTex 05	4.74
Deep Maize 06 F	3.58
Deep Maize 06 SF	n/a
MinneTAC 06	-2.70
PhantAgent 06	4.15
TacTex 06	5.85

Table 6: Eleven TAC/SCM agents constituting the SCM_{↓3} 2005 & SCM_{↓3} 2006 strategy sets, along with their results from their respective year’s final tournament round (in \$M).

Using the resultant SCM_{↓3} 2005 and SCM_{↓3} 2006 games we compare TABU and BFS on the two measures of performance. The SCM_{↓3} 2005 has no PSNE and we therefore use this game to compare the algorithms on the search required to approximate equilibria at varying levels of ϵ . Starting from each of the 56 distinct profiles we run each algorithm and report the mean minimum confirmed ϵ given by the percentage of space searched. The result is shown in Figure 5. On average, BFS confirms low ϵ profiles sooner. In some cases, TABU fails to ever confirm the lowest ϵ profile in the game, resulting in a higher mean estimate when considering

60% or more of the space searched.




Figure 5: Mean low ϵ for $SCM_{\downarrow 3}$ 2005.

The second game, $SCM_{\downarrow 3}$ 2006, does have a PSNE. Starting from each of the now 35 distinct profiles we run each algorithm and report the mean search required to confirm the PSNE. The result is shown in Figure 6. The mean percentage of space searched is 54.0% for BFS and 56.6% for TABU. For TABU there were four cases where the algorithm failed to confirm the PSNE.

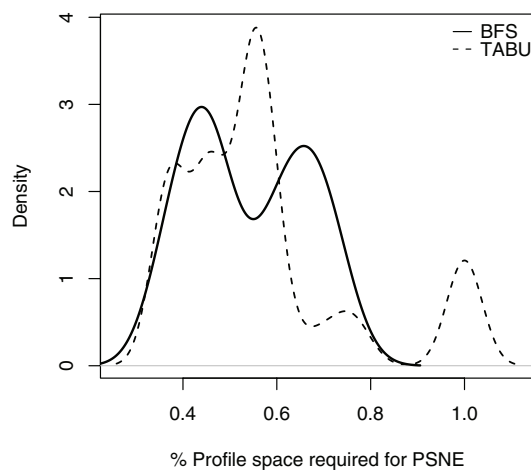


Figure 6: Kernel density plot showing the space required for finding a PSNE in the $SCM_{\downarrow 3}$ 2006 game.

The analysis for the individual $SCM_{\downarrow 3}$ games considered by Jordan, Kiekintveld, & Wellman (2007) totaled 3210 val-

idated game instances each requiring 7 CPU hours to perform. A major component of their analysis was identifying approximate or PNSE which give a background context for comparison (NE-response ranking) and for which a best response can be designed. Using BFS, the minimum ϵ profile could have been determined using less than 60% of the compute cycles required for the full game analysis.

Discussion

We have presented empirical results comparing BFS and TABU on classes of algorithms varying in structure from weakly structured to strongly structured games. In all cases, the performance comparison of BFS and TABU on the space exploration measure shows both are essentially equivalent. An important attribute of the BFS algorithm is that it will confirm all available profiles eventually, whereas TABU will not necessarily do so. This is important not only in the case where no PSNE exists, but also when we wish to analyze low ϵ profiles when designing a best response. For many games of interest it may be that the second proposed measure may provide a more relevant statistic. We have shown empirically that BFS outperforms TABU on games with little structure with this measure.

We have shown that in classes of games which there is underlying structure (in this case potential and LEGs) both algorithms converge quickly. If an empirical game is believed to possess characteristics of a known class of games, for instance LEGs, it may be reasonable to model the empirical game as an approximate LEG. In this case we could use other sampling methods to approximate the underlying (potential) functions. A comparison of BFS to such an approach would provide a good understanding of the benefits of these algorithms in the case where we have added intuition about the structure of the game.

Acknowledgments

This work was supported in part by the National Science Foundation grant 0414710 and the National Science Foundation IGERT grant 0114368.

References

- Arunachalam, R., and Sadeh, N. M. 2005. The supply chain trading agent competition. *Electronic Commerce Research and Applications* 4:63–81.
- Eriksson, J.; Finne, N.; and Janson, S. 2006. Evolution of a supply chain management game for the trading agent competition. *AI Communications* 19:1–12.
- Jordan, P. R.; Kiekintveld, C.; and Wellman, M. P. 2007. Empirical game-theoretic analysis of the TAC supply chain game. In *Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems*.
- Kiekintveld, C.; Wellman, M. P.; and Singh, S. 2006. Empirical game-theoretic analysis of Chaturanga. *AAMAS-06 Workshop on Game-Theoretic and Decision-Theoretic Agents*.
- Leyton-Brown, K., and Tennenholtz, M. 2003. Local-effect games. In *Eighteenth International Joint Conference on Artificial Intelligence*, 772–780.

- Monderer, D., and Shapley, L. S. 1996. Potential games. *Games and Economic Behavior* 14:124–143.
- Nudelman, E.; Wortman, J.; Shoham, Y.; and Leyton-Brown, K. 2005. Run the GAMUT: A comprehensive approach to evaluating game-theoretic algorithms. In *Third International Joint Conference on Autonomous Agents and Multiagent Systems*, 880–887.
- Reeves, D. M.; Wellman, M. P.; MacKie-Mason, J. K.; and Osepashvili, A. 2005. Exploring bidding strategies for market-based scheduling. *Decision Support Systems* 39:67–85.
- Reeves, D. M. 2005. *Generating Trading Agent Strategies: Analytic and Empirical Methods for Infinite and Large Games*. Ph.D. Dissertation, University of Michigan.
- Rosenthal, R. W. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2:65–67.
- Russell, S., and Wefald, E. 1991. Principles of metareasoning. *Artificial Intelligence* 49:361–395.
- Sureka, A., and Wurman, P. R. 2005. Using tabu best-response search to find pure strategy Nash equilibria in normal form games. In *Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, 1023–1029.
- Vorobeychik, Y.; Kiekintveld, C.; and Wellman, M. P. 2006. Empirical mechanism design: Methods, with application to a supply-chain scenario. In *Seventh ACM Conference on Electronic Commerce*, 306–315.
- Walsh, W.; Parkes, D.; and Das, R. 2003. Choosing samples to compute heuristic-strategy Nash equilibrium. In *Fifth Workshop on Agent-Mediated Electronic Commerce*.
- Wellman, M. P.; Reeves, D. M.; Lochner, K. M.; Chen, S.-F.; and Suri, R. 2005. Approximate strategic reasoning through hierarchical reduction of large symmetric games. *Twentieth National Conference on Artificial Intelligence* 502–508.
- Wellman, M. P. 2006. Methods for empirical game-theoretic analysis (extended abstract). In *Twenty-First National Conference on Artificial Intelligence*, 1552–1555.