

Scaling Simulation-Based Game Analysis through Deviation-Preserving Reduction

Bryce Wiedenbeck and Michael P. Wellman
University of Michigan
{btwied,wellman}@umich.edu

ABSTRACT

Multiagent simulation extends the reach of game-theoretic analysis to scenarios where payoff functions can be computed from implemented agent strategies. However this approach is limited by the exponential growth in game size relative to the number of agents. Player reductions allow us to construct games with a small number of players that approximate very large symmetric games. We introduce *deviation-preserving reduction*, which generalizes and improves on existing methods by combining sensitivity to unilateral deviation with granular subsampling of the profile space. We evaluate our method on several classes of random games and show that deviation-preserving reduction performs better than prior methods at approximating full-game equilibria.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics

General Terms

Algorithms, Economics

Keywords

empirical game theory, simulation-based game theory, game reduction

1. INTRODUCTION

Game-theoretic analysis plays an increasingly prominent role in research on understanding and designing multiagent systems. Agent-based simulation offers the potential to increase the scope of applicability for game theory, beyond those game scenarios that can be described straightforwardly and solved analytically. In the simulation-based approach, rather than directly express all payoffs for a game, the analyst describes an environment procedurally and then computes payoffs by simulation of agent interactions in that environment.

Simulation enables analysis of many rich strategic environments, but determining payoffs for a large game in this

way may be prohibitively expensive. Straightforward estimation of a payoff function requires simulation of every possible combination, or *profile*, of agent strategies. If the environment is stochastic, then many simulation runs may be necessary to obtain a reasonable estimate of even a single profile. For multiagent interactions that extend over time, or are otherwise complex, the computational cost of simulation may severely limit the number of profiles—and therefore the size of the game—that can be considered in such an analysis.

We focus for most of this paper on *symmetric games*, in which all agents have the same set of available strategies and payoffs depend only on the number of agents playing each strategy, not on the specific identities of those agents. Formally, a symmetric game is a tuple $\Gamma = (N, S, u)$, where N is the number of agents, S is the set of strategies available to all agents, and the utility function $u(s, \vec{s})$ gives the payoff to any agent playing strategy s in profile \vec{s} . To conduct a complete analysis of Γ , we require that u specifies payoffs for all possible profiles. A symmetric game with N agents and $|S|$ strategies contains $\binom{N+|S|-1}{N}$ profiles.¹ For a sense of how great a burden this imposes, consider that a symmetric game with 15 agents and 15 strategies contains over 77 million profiles, so if estimating a profile's payoff through simulation required one second, constructing the full game would take more than two years.

We seek to combat this exponential growth using a technique broadly known as *player reduction*. Player reductions approximate games with many agents by constructing smaller games that aggregate over those agents in some way. Equilibria of the reduced game can then be viewed as approximate equilibria of the full game.

As an example, consider trading in continuous double auctions (CDAs)—a problem of agent strategy that has been extensively investigated through simulation. We review the coverage of several studies that employed simulation to estimate payoff functions for purposes of game-theoretic or evolutionary analysis. In the first empirical game analysis of CDA strategy, Walsh et al. [15] analyzed a 20-player game with three strategies. The 231 distinct profiles were within their simulation budget, whereas adding just one

¹To see this, note that we can describe a profile in terms how many agents play each strategy. Suppose an ordering of strategies, and consider a representation that indicates players by one symbol (.) and partitions by another (|). For instance, with $S = \{s_1, s_2, s_3, s_4\}$ and $N = 6$, the profile $\dots|..|.$ has three agents playing s_1 , two s_2 , and one s_4 . The representation contains $N + |S| - 1$ total symbols, and the choice of which N of them to make players (or equivalently, choice of partitions) uniquely defines a profile.

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), June, 4–8, 2012, Valencia, Spain.

Copyright © 2012, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

more strategy would have entailed estimating 1540 more. Vytelingum et al. [14] likewise considered a 20-player game, though one that imposed symmetry only within the subgroups of 10 buyers and 10 sellers. Their study also compared three strategies, but for tractability and to facilitate visualization of their evolutionary traces, they limited analysis to two strategies at a time, which requires 121 profiles for each strategy pair and scenario combination.² Phelps et al. [8] covered up to four strategies, in a 12-agent simulation employing another form of double auction mechanism (455 profiles). The study of Tesouro and Bredin [11] considered as many as 44 trading agents with three different strategies, but their analysis evaluated only profiles where agents were evenly divided across two strategies. This selection is similar to a two-player reduction according to the hierarchical method, as discussed below. Tesouro and Das [12] covered five strategies for a 20-agent scenario, this time with a mix of evenly-divided profiles and profiles where only one agent deviates from a homogeneous profile. As we see below, this is suggestive of the deviation-preserving reduction method we introduce here. In by far the most comprehensive CDA simulation study to date, Schwartzman and Wellman [10] systematically evaluated 14 strategies in a 16-agent scenario. This was rendered feasible only by virtue of their reduction to a four-player game, comprising 2380 profiles as opposed to 68 million in the unreduced game. Overall, we see that many-agent simulation studies either adopt player reductions, or make do with very narrow strategy exploration.

In this paper, we propose and study *deviation-preserving reduction*, which renders reduced-game equilibria more informative with respect to the full game. We start in Section 2 by reviewing existing methods for player reduction. Section 3 introduces deviation-preserving reduction, and explains how our new method is designed to combine the best aspects of its predecessors. Section 4 evaluates the reductions, and Section 5 shows how both our reduction and previous ones can be extended to games that are symmetric only with respect to a partition of players into roles.

2. BACKGROUND: PLAYER REDUCTION

Two methods for player reduction have been proposed in the literature: *hierarchical reduction* [16], and *twins reduction* [4]. Both methods are defined with respect to symmetric games. They define a subset of the profiles in the given (*full*) game, and map the payoffs of these profiles to a payoff function defined over a game with fewer players (the *reduced game*). Analyses of the reduced game are then interpreted as approximately applying to the full game.

2.1 Hierarchical Reduction

Of the two existing methods, hierarchical reduction is the more extensively used [1, 5, 10]. Hierarchical reduction works by grouping agents into coalitions that are constrained to act together. One player in the reduced game selects an action to be played by all agents in a coalition, and receives the payoff to any agent playing that strategy. To capture this formally, we introduce the following notation: a strategy profile $\vec{s} = \langle c_1 \times s_1, \dots, c_{|S|} \times s_{|S|} \rangle$ of game Γ consists of strategies $s_i \in S$ and integer counts $c_i \geq 0$ for each strategy such that $\sum_{i=1}^{|S|} c_i = N$. When $c_i = 0$, we may omit it from

²These same authors in earlier work [13] evaluated a 20-player three-strategy CDA game.

the expression.

The hierarchical reduction of Γ to $n < N$ players is defined as $HR_n(\Gamma) = (n, S, u^{HR})$, where

$$u^{HR} (s, \langle c_1 \times s_1, \dots, c_{|S|} \times s_{|S|} \rangle) = u \left(s, \left\langle \frac{N}{n} c_1 \times s_1, \dots, \frac{N}{n} c_{|S|} \times s_{|S|} \right\rangle \right).$$

This definition follows previous applications of hierarchical reduction in assuming that N is an integer multiple of n . In our evaluation we employ a generalized version (described in Section 4.1) that allows reduction to numbers of players that do not evenly divide the number of agents in the full game.

For illustration, consider a full game with $N = 25$ agents, and a hierarchical reduction to $n = 5$ players. The action of each reduced-game player is played by five agents in the full game, so the reduced-game profile $\langle 2 \times s_1, 1 \times s_2, 2 \times s_3 \rangle$ corresponds to the full-game profile $\langle 10 \times s_1, 5 \times s_2, 10 \times s_3 \rangle$.

The main idea behind hierarchical reduction is that though the payoff to a particular strategy generally varies with the number of agents that play each strategy, it often can be expected to do so smoothly. Kearns and Mansour [6] formalize a related condition called *bounded influence* to define a class of compactly representable and solvable games. Whereas it is easy to construct games that violate this assumption, in many natural symmetric games, the payoffs are smooth in this way.

However, $HR_n(\Gamma)$ lacks crucial information relevant to Nash equilibria of Γ . In a Nash equilibrium of Γ , no individual agent can gain by deviating to another strategy, but the hierarchical reduction contains no information about unilateral deviations. In an equilibrium of $HR_n(\Gamma)$, no $\frac{N}{n}$ -agent coalition can gain by all deviating to the same strategy, but there are many cases in which these conditions differ substantially.

Consider a network formation game [3] in which agents create links to one another. Agents gain from being in a connected network, but incur a cost for each link they create. In such a game we can envision a full-game equilibrium that is not an equilibrium of the reduced game, as well as a reduced-game equilibrium that is not an equilibrium of the full game. If network effects are large, but cost of creating links is high, there could be an equilibrium of the full game where agents create no links. However, if several players are allowed to deviate together they may create a sufficiently dense network to overcome the link-creation cost: this would be a beneficial deviation in the reduced game, so the full-game equilibrium would not be found. Under different parameters, there may be a spurious equilibrium in the reduced game where all players contribute links to the network, and no player can gain by deviating because if all agents represented by one reduced-game player changed strategies simultaneously, the network would collapse. On the other hand, a unilaterally deviating agent in the full game might have a much smaller impact and still receive the network benefits while avoiding the link creation cost.

2.2 Twins Reduction

The natural solution to this problem is to incorporate information about the value of unilateral agent deviations into the payoffs of the reduced game. Ficici et al. [4] propose a method called *twins reduction* that takes a first step in this

direction. The twins reduction of a symmetric game³ is a 2-player game, $TR(\Gamma) = (2, S, u^{TR})$, where each player views itself as controlling one agent in the full game, and the opponent as controlling all remaining agents:

$$u^{TR}(s, \langle 1 \times s, 1 \times s' \rangle) = u(s, \langle 1 \times s, (N-1) \times s' \rangle).$$

Note that the payoffs for the two strategies in a twins reduction profile $\langle 1 \times s, 1 \times s' \rangle$, $s \neq s'$, correspond to two different profiles in the full game:

- $\langle 1 \times s, (N-1) \times s' \rangle$ and
- $\langle (N-1) \times s, 1 \times s' \rangle$,

but that the reduced game is still symmetric.

Ficici et al. [4] advocate constructing twins reduction games not by explicitly simulating these full-game profiles, but by sampling random profiles from the full game and determining the payoffs by linear regression on the number of agents playing each strategy. We refer to this approach as *TR-R*, where the second “R” stands for “regression”. We consider the direct simulation approach a more appropriate benchmark, but evaluate both methods in our experiments.

The advantage of the twins reduction is that it captures information about individual agents’ incentives to deviate. Its major disadvantage is that it is limited to two players, and can therefore give only an extremely coarse-grained view of the game. In general, a reduced-game representation will have difficulty capturing equilibria of the full game that have support size (number of distinct strategies played with positive probability) larger than n , the reduced number of players, as no profiles of the reduced game capture the interaction of all strategies in the support set. Since the twins reduction ($n = 2$) never contains profiles where more than two strategies are played, it is particularly restrained by this limitation.

3. DEVIATION-PRESERVING REDUCTION

We propose a new game reduction method that combines the sensitivity to unilateral deviation afforded by twins reduction with the profile-space granularity of hierarchical reduction. We call this method *deviation-preserving reduction*. In a deviation-preserving reduction game, each player views itself as controlling a single agent in the full game, but views the profile of opponent strategies in the reduced game as an aggregation of all other agents in the full game. Formally, $DPR_n(\Gamma) = (n, S, u^{DPR})$, where

$$u^{DPR}(s, \langle c_1 \times s_1, \dots, c_s \times s, \dots \rangle) = u\left(s, \left\langle \frac{N-1}{n-1} c_1 \times s_1, \dots, \left[\frac{N-1}{n-1} (c_s - 1) + 1 \right] \times s, \dots \right\rangle\right).$$

In a hierarchical reduction, the proportion of agents playing each strategy is the same in the full and reduced games. Under deviation-preserving reduction, analogously, the proportion of *opponents* playing a strategy in the full and reduced games is the same from each player’s perspective. And as in a twins reduction, each player in a deviation-preserving

³The original definition [4] applies to a somewhat broader class: role-symmetric games with identical strategies. In Section 5 we describe how to extend player reduction, including twins reduction, to the entire class of role-symmetric games.

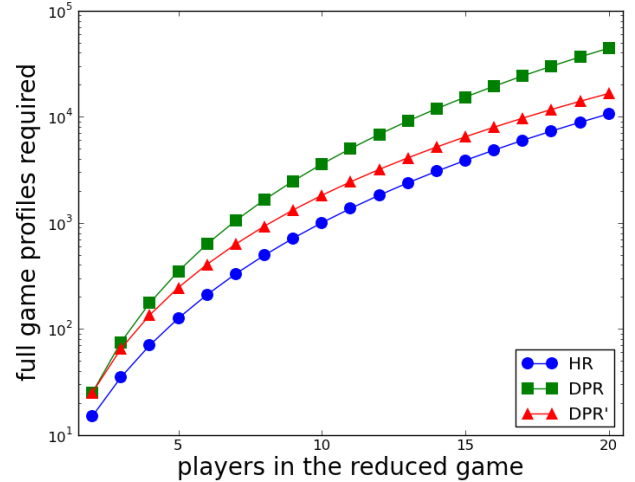


Figure 1: Number of full-game profiles required to construct reduced games (log scale), for $|S| = 5$.

reduction game is sensitive to the payoffs of exactly one agent in the full game. As a consequence of this sensitivity to single agents, the deviation-preserving reduction game can identify exact symmetric pure strategy equilibria of the full game if they exist.

PROPOSITION 1. *A profile $\langle n \times s \rangle$ is a Nash equilibrium of $DPR_n(\Gamma)$ if and only if the profile $\langle N \times s \rangle$ is a Nash equilibrium of Γ .*

PROOF. The profile $\langle n \times s \rangle$ is a NE when $u^{DPR}(\langle n \times s \rangle) \geq u^{DPR}(s, \langle (n-1) \times s, 1 \times s' \rangle)$ for all $s' \in S$. This is the case exactly when $u(s, \langle N \times s \rangle) \geq u(s, \langle (N-1) \times s, 1 \times s' \rangle)$ for all $s' \in S$. \square

This property also holds for twins reduction games, because the deviation preserving reduction is a strict generalization of the directed-sampling twins reduction: $TR(\Gamma) = DPR_2(\Gamma)$.

To construct each profile’s payoffs in a deviation-preserving reduction game, several profiles from the full game must be simulated. Returning to the example of a 25-agent full game and a 5-player reduced game, the profile $\langle 2 \times s_1, 1 \times s_2, 2 \times s_3 \rangle$ in the deviation-preserving reduction game employs payoff values from several profiles in the full game:

$$\begin{aligned} u^{DPR}(s_1, \langle 2 \times s_1, 1 \times s_2, 2 \times s_3 \rangle) &= u(s_1, \langle 7 \times s_1, 6 \times s_2, 12 \times s_3 \rangle) \\ u^{DPR}(s_2, \langle 2 \times s_1, 1 \times s_2, 2 \times s_3 \rangle) &= u(s_2, \langle 12 \times s_1, 1 \times s_2, 12 \times s_3 \rangle) \\ u^{DPR}(s_3, \langle 2 \times s_1, 1 \times s_2, 2 \times s_3 \rangle) &= u(s_3, \langle 12 \times s_1, 6 \times s_2, 7 \times s_3 \rangle) \end{aligned}$$

Note that we again assume divisibility: in this case, $n-1$ has to divide $N-1$ for the aggregation of opponents to be precise. As with hierarchical reduction, we can extend the definition to reduced games with any number of players, as described in Section 4.1. We quantify the number of profile simulations required for deviation-preserving reduction in the following proposition.

PROPOSITION 2. *Constructing $DPR_n(\Gamma = (N, S, u))$ requires simulating $|S| \binom{n+|S|-2}{n-1}$ full-game profiles.*

PROOF. In each profile of the deviation-preserving reduction game, $n - 1$ of the players each control $\frac{N-1}{n-1}$ full-game agents. The set of all such profiles can be viewed as an $(n - 1)$ -player symmetric game, so we know that there are $\binom{n+|S|-2}{n-1}$ of them. Each of these profiles must be paired with each $s \in S$, so $|S|\binom{n+|S|-2}{n-1}$ profiles must be simulated. \square

Proposition 2 shows that constructing $DPR_n(\Gamma)$ requires simulating strictly more profiles than $HR_n(\Gamma)$, but by a factor of at most $|S|$. As we show in Section 4, the extra profiles comprising this constant factor can contribute to significantly improved accuracy. Even so, we would like to minimize the number of simulations required when possible, and therefore also consider a variant of the deviation-preserving reduction, which we call DPR' .

The idea behind DPR' is that many of the profiles simulated to construct a deviation-preserving reduction game are quite similar. For example, in the 5-player deviation-preserving reduction of the 25-agent game, the payoff to strategy s_1 in the full-game profile $\vec{s}_a = \langle 7 \times s_1, 6 \times s_2, 12 \times s_3 \rangle$ is employed in the reduced-game profile $\langle 2 \times s_1, 1 \times s_2, 2 \times s_3 \rangle$. The payoff for s_2 in reduced-game profile $\langle 1 \times s_1, 2 \times s_2, 2 \times s_3 \rangle$ is derived from $\vec{s}_b = \langle 6 \times s_1, 7 \times s_2, 12 \times s_3 \rangle$, which differs from \vec{s}_a only in that a single agent has switched from s_1 to s_2 , both of which are played by many other agents. If we believe our assumption—inherited from hierarchical reduction—that payoffs vary smoothly in the number of agents playing each strategy, we should expect the payoffs to strategy s_1 in profiles \vec{s}_a and \vec{s}_b to be very similar (likewise for s_2), suggesting that we could get away with simulating only one of the two.

Formally, $DPR'(\Gamma) = (n, S, u^{DPR'})$, where $u^{DPR'}$ is defined as follows. Let $\vec{s} = \langle c_{\min} \times s_{\min}, \dots, c_s \times s, \dots \rangle$, where s_{\min} is the first strategy played by at least one agent. If $c_s = 1$ or $c_s = c_{\min}$, then $u^{DPR'}(s, \vec{s}) = u^{DPR}(s, \vec{s})$. Otherwise, $u^{DPR'}(s, \vec{s})$ is given by

$$u(s, \left\langle \left[\frac{N-1}{n-1} c_{\min} + 1 \right] \times s_{\min}, \dots, \frac{N-1}{n-1} (c_s - 1) \times s, \dots \right\rangle).$$

The result is that when DPR would prescribe simulation of several profiles that differ only by deviation of a single agent (and no strategy is played by only one agent), DPR' requires that only one be simulated. From among these profiles, DPR' selects the one in which the lowest-numbered strategy by which they differ is played most. In the example above, payoffs

- $u^{DPR'}(s_1, \langle 2 \times s_1, 1 \times s_2, 2 \times s_3 \rangle)$ and
- $u^{DPR'}(s_2, \langle 1 \times s_1, 2 \times s_2, 2 \times s_3 \rangle)$

both come from full-game profile $\langle 7 \times s_1, 6 \times s_2, 12 \times s_3 \rangle$. The savings in terms of profiles sampled are illustrated in Figure 1. In that graph, the curve for DPR' follows the formula $|\Gamma| = |S|\binom{n+|S|-2}{n-1} - (n-2)\binom{n+|S|-3}{n-1}$. DPR' always requires more full-game profiles than HR , and fewer than DPR , except when $n = 2$, where both are equivalent to TR .

4. EMPIRICAL EVALUATION

The goal of a player reduction is to replace a full game that is too large to effectively analyze with a more manageable reduced game. To compare reduction methods, we therefore need to evaluate how well analysis performed on a reduced

game translates back to the full game. This presents a problem for evaluation, in that full games of interest are too big to effectively analyze. For example, in the simulated credit network games discussed below, we construct 12-agent, 6-strategy full games; we would like to analyze reductions of 60-agent games, but even with just six strategies, the full game would consist of 8,259,888 profiles. We therefore compromise by reducing several types of medium-sized games to very small ones. If one reduction consistently performs better in such cases, we take it as an indication that the same will hold for reductions of very large games.

4.1 Regret of Reduced Game Equilibria

Numerous methods for analyzing games exist, but the most important is finding Nash equilibria. Because player reductions work with symmetric games, we evaluate them primarily by how well symmetric mixed strategy Nash equilibria computed in the reduced game approximate symmetric mixed strategy equilibria of the full game. Our primary measure for the quality of reduced-game equilibria is *regret*. The regret $\epsilon(\vec{\sigma})$ of a symmetric mixed strategy profile $\vec{\sigma}$, in which all players play mixed strategy σ is the maximum gain any player could achieve by deviating to a pure strategy:

$$\epsilon(\vec{\sigma}) = \max_{s \in S} u(s, \vec{\sigma}_{-i}),$$

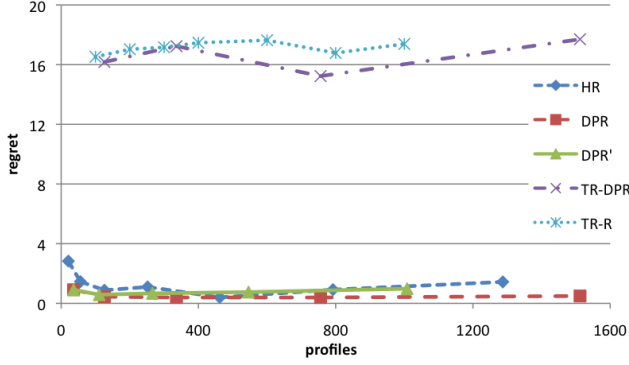
where $u(s, \vec{\sigma}_{-i})$ is the expected payoff to a player playing s when all others play σ . A Nash equilibrium has zero regret, but a symmetric mixed profile $\vec{\sigma}$ that is an equilibrium of the reduced game will generally have positive regret with respect to the full game. Such a $\vec{\sigma}$ can be viewed as an *approximate*, or $\epsilon(\vec{\sigma})$ -Nash equilibrium of the full game, where the lower the regret, the better the approximation.

However, we cannot simply compare the regret of equilibria from k -player reduced games under each method. The first problem is that the number of players in a twins reduction game is not scalable. Moreover, since the goal of player reduction is to simulate fewer profiles, the relevant comparison is not the number of players in the reduced game, but the number of profiles required to construct it, and DPR_k always requires sampling strictly more profiles than HR_k . For example, in the 12-agent 6-strategy game instances below, $|DPR_3| = 126 = |HR_4|$. Because the directed-simulation twins reduction always requires $\frac{|S|^2(|S|-1)}{2}$ profiles, we compare to TR by addressing the question of whether lower regret can be achieved by a method that samples more profiles. Twins reduction with regression can use any set of profiles; in our experiments we varied the size of the set over a range similar to that required to construct the various reduced games.


Hierarchical reduction as defined by Wellman et al. [16] requires that the number of reduced-game players n divide the number of full-game agents N . By analogy, in our definition of deviation preserving reduction above, we assume that $n - 1$ divides $N - 1$. In our experiments, we perform reductions of both varieties where these conditions do not hold. We extend the definition of HR_n to allow indivisibility as follows.

$$u^{HR}(s, \langle c_1 \times s_1, \dots, c_{|S|} \times s_{|S|} \rangle) = u(s, \left\langle \left\lfloor \frac{N}{n} c_1 + 1 \right\rfloor \times s_1, \dots, \left\lfloor \frac{N}{n} c_{j+1} \right\rfloor \times s_{j+1}, \dots \right\rangle),$$

where $j = N - \sum_i \lfloor \frac{N}{n} c_i \rfloor$. That is, the number of opponents



(a) all reduction types



(b) rescaled to exclude regression-based twins reductions

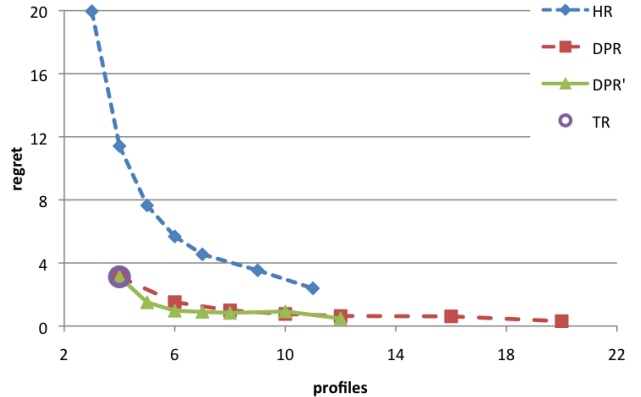
Figure 2: Average full-game regret of reduced-game equilibria in local effect games. $N = 12$, $|S| = 6$, $2 \leq n \leq 8$.

playing strategy s_i is the integral part of $\frac{N}{n}c_i$, with extra player slots allocated one each to strategies with lower indices. For example, when we construct HR_5 of a game with 12 players, the reduced-game profile $(1 \times s_1, 3 \times s_2, 1 \times s_3)$ corresponds to the full-game profile $(3 \times s_1, 7 \times s_2, 2 \times s_3)$. We extend the definition of DPR_n to handle indivisibility in a similar manner.

We evaluate the reductions using two classes of random games: congestion games [9], in which agents select a fixed-size subset of available facilities and payoffs are decreasing in the number of agents choosing a facility; and local effect games [7], which have a graph over actions and each action’s payoff is a function of the number of agents choosing it and adjacent actions. We randomly varied the payoff function parameters of these games to create 250 game instances for each test described below. We also evaluate on one simulated game class, based on a scenario of *credit network formation* [2]. In the model of credit networks employed in this scenario, directed links represent credit issued to other agents: agents wish to transact with one-another, but issuing credit bears risk in that debtors may default. In the credit network formation game, payoffs are determined by simulating a sequence of transactions and defaults on the network induced by agent strategies. We sampled each profile of a 12-agent, 6-strategy credit network game 100 times, and randomly recombined these samples to create 250 game instances.

The first finding of note is that twins reduction performs very poorly with linear regression. The top line in Figure 2a shows the regret of equilibria found in $TR-R$ games with random sampling of profiles, which is an order of magnitude worse than the HR , TR , DPR , and DPR' . Two observations led us to try the method labeled $TR-DPR$: first, that sampling profiles according to uniform agent play leads to a very low likelihood of observing payoffs for profiles where most agents play the same strategy, and these are exactly the profiles whose payoffs the regression estimates. Second, simulating all the profiles for DPR or DPR' makes available a substantial amount of payoff data that goes unused in constructing the reduced game. We therefore thought to try using all of the profiles simulated for the deviation-preserving reduction as input to the linear regression of $TR-R$. As is clear from Figure 2a, this improves very little on random sampling.

In retrospect, it is not particularly surprising that approximating payoffs by linear regression performs so poorly: all of our example games and most games requiring simulation have nonlinear payoffs. A better regression model could potentially alleviate this problem, but choosing one requires knowledge of the game’s payoff function that may not be available when payoffs are determined by simulation. We also ran $TR-R$ and $TR-DPR$ on each of the other game classes, but the results are similarly poor, and are excluded from subsequent figures.

Figure 3: Full-game regret of reduced-game equilibria in congestion games. $N = 100$, $|S| = 2$, $2 \leq n \leq 10$.

Figures 2b, 3, and 4 show that deviation-preserving reduction outperforms hierarchical reduction and twins reduction in a wide variety of settings. In 12-agent, 6-strategy local effect games, DPR is clearly better than HR , but the comparison to DPR' is less conclusive. We were surprised to find that hierarchical reduction would perform worse with increased reduced-game size, which corresponds to increased abstraction granularity. We note, however, that the 5, 7, and 8-player reduced games where HR performs poorly are exactly the cases where n does not divide $N = 12$. This also leads us to observe that because 11 is prime, the deviation-preserving reduction never has the advantage $n - 1$ dividing $N - 1$, and yet consistently performs well. The results from 12-agent, 6-strategy congestion games (not shown) are broadly similar.




Figure 4: Average full-game regret of reduced-game equilibria in credit network games. $N = 12$, $|S| = 6$, $2 \leq n \leq 8$.

In an attempt to get at the effect of very substantial player reductions, we created 100-agent, 2-strategy congestion games. The results in Figure 3 show clear separation between hierarchical reduction and both variants of deviation-preserving reduction, suggesting that as the number of players grows, the relative difference between DPR and DPR' may be smaller. Results for the 12-agent, 6-strategy credit network game appear in Figure 4. Here again, DPR and DPR' perform similarly, and better than HR .

Across all game classes and sizes examined (including those not shown) deviation-preserving reduction of any given size outperforms the hierarchical reduction with at least as many profiles that is closest in size. This means that for any size hierarchical reduction, there exists a better deviation-preserving reduction that requires simulating fewer profiles. Virtually all of these differences are significant at $p < 0.05$; the only exceptions are 4-player DPR versus 6-player HR in the 12-player congestion game (Figure 3) and 12-player credit network game (Figure 4). In addition, DPR_3 outperforms TR across all game classes; the difference is significant at $p < 0.05$ in all cases except the credit network game. The difference between DPR_4 and TR is significant in all cases.

4.2 Comparison to Full-Game Equilibria

We also compared reduced-game equilibria under HR and DPR to equilibria from 12-player, 6-strategy full games using two metrics: similarity of support sets, and L^2 distance between distributions. Table 1 shows the number of strategies by which the support sets of full and reduced-game equilibria differ. Here, we consider a strategy to be in the support of a symmetric ϵ -Nash equilibrium if it is played with probability 0.01 or greater. In nearly all cases support sets of DPR_n match those of full-game equilibria significantly ($p < 0.05$) better than both HR_n and HR_{n+1} . In addition, for congestion games and local effect games, $DPR_{>2}$ significantly outperforms TR , whereas in credit network games, there is no significant difference between TR and DPR .

Table 2 presents a similar message, but in terms of the L^2 distances between the mixed strategy distributions in full and reduced-game equilibria. Again, DPR is significantly better than HR and TR for congestion and local effect games, while performing similarly on credit network games. As in Section 4.1, in these experiments, we compute

	credit network		congestion		local effect	
n	HR	DPR	HR	DPR	HR	DPR
2	3.72	1.49†	1.68	0.39†	2.72	0.45†
3	3.72	1.64†	0.99*	0.17†*	1.04*	0.20†*
4	3.72	1.60†	1.05	0.10†*	0.98	0.12†*
5	1.98*	1.54†	1.10	0.08†	1.01	0.09†
6	1.19†*	1.39	0.98	0.05†	0.85*	0.08†

Table 1: Reduced versus full-game NE support set difference. * indicates significant difference between n and $n - 1$; † indicates significant difference between HR_n and DPR_n .

	credit network		congestion		local effect	
n	HR	DPR	HR	DPR	HR	DPR
2	0.713	0.703	0.435	0.069†	0.503	0.128†
3	0.764*	0.690†	0.141*	0.039†*	0.154*	0.064†*
4	0.860	0.640†	0.117*	0.022†*	0.117*	0.037†*
5	0.643*	0.641	0.141	0.019†*	0.144	0.027†*
6	0.467†*	0.564*	0.099*	0.018†	0.088*	0.026†

Table 2: Reduced versus full-game NE distribution L^2 distance. * indicates significant difference between n and $n - 1$; † indicates significant difference between HR_n and DPR_n .

one symmetric mixed-strategy Nash equilibrium per game by running replicator dynamics initialized to the uniform mixture.

4.3 Dominated Strategies

Another useful operation in the analysis of simulation-based games is to check for dominated strategies. A dominated strategy is one that no agent should ever play because there is an alternative strategy that is at least as good in response to any profile of opponent strategies. We ran experiments on 12-agent, 6-strategy congestion and credit network games (250 each), comparing the set of strategies that remain after iterated elimination of strictly dominated strategies in the full game against those that remain in 2, 4, and 6-player reduced games. We observed that DPR and DPR' produced very similar results, and that both improved over hierarchical and twins reduction. Figures 5 and 6 show histograms of the number of strategies eliminated in reduced games but not eliminated in full games.

In congestion games (Figure 5), twins reduction and both forms of deviation-preserving outperform hierarchical reduction, eliminating fewer strategies in the reduced game that survive in the full game, even when hierarchical reduction samples vastly more profiles. These congestion games often exhibit dominated strategies in the full game, but we almost never observed strategies surviving in reduced games that are dominated in the full game.

In credit network games (Figure 6), no strategies are dominated in the full game, but in the twins reduction game, many strategies are eliminated. Moving to DPR_4 or DPR'_4 solves this problem almost entirely. These experiments also confirm that for all reduction types, increasing the number of players in the reduced game reduces the number of strategies erroneously found to be dominated.

5. ROLE-SYMMETRIC GAMES

We can smoothly relax the constraint that games be fully




Figure 5: Histograms showing the number of strategies surviving iterated elimination of dominated strategies in full but not reduced congestion games. $N = 12$, $|S| = 6$, 250 random games. $TR \equiv DPR_2$ outperforms HR , sampling far fewer profiles.




Figure 6: Histograms showing the number of strategies surviving iterated elimination of dominated strategies in full but not reduced credit network games. $N = 12$, $|S| = 6$, 250 sample games. DPR'_4 avoids the aggressive elimination occurring in TR .

symmetric by assigning agents to roles, and enforcing symmetry only within these roles. Across roles, agents' strategy sets and payoffs can differ, but within a role, they are symmetric. Formally, a *role-symmetric game* is a tuple $\Gamma = (\{N_i\}, \{S_i\}, u)$, where the number of agents with role i is N_i , and agents with role i have strategy set S_i . Role-symmetric games provide a natural model for many settings where agents can be partitioned into meaningful categories, such as buyers and sellers in a market, or attackers and defenders in a security game. Role symmetry imposes no loss of generality on normal-form games, spanning the spectrum from complete asymmetry (each player has its own role) to full symmetry (a single role for everyone).

All of the player reduction methods discussed here can be straightforwardly extended to role-symmetric games. Consider for example the 20-agent continuous double auction study of Vytelingum et al. [14] with $N_1 = 10$ buyers and $N_2 = 10$ sellers. Instead of choosing n , the number of players in the reduced game, we must choose each $\{n_i\}$, the number of players with each role in the reduced game.

To perform a hierarchical reduction, a natural choice would be $n_1 = n_2 = 2$. This would involve simulating all profiles where 0, 5, or 10 agents play each buyer strategy, and a multiple of five agents likewise play each seller strategy.

With twins reduction, there are two players per role. Each

player views itself as controlling a single agent, and the other player with the same role as controlling nine agents. It views the two other-role players as each representing half the ten agents with that role, so in the reduced-game profile

$$\langle 1 \times s_{1.1}, 1 \times s_{1.2}, 1 \times s_{2.1}, 1 \times s_{2.2} \rangle,$$

the payoff to buyer 1, who plays $s_{1.1}$, comes from full-game profile

$$\langle 1 \times s_{1.1}, 9 \times s_{1.2}, 5 \times s_{2.1}, 5 \times s_{2.2} \rangle.$$

The deviation-preserving reduction extends the twins reduction to more than two reduced-game players per role, maintaining the view that a reduced-game player controls a single agent, while the other players with the same role aggregate over *the rest* of the agents with that role, and players with another role aggregate over *all* agents with their role. With either hierarchical reduction or deviation-preserving reduction, it would be possible to choose $n_i \neq n_j$ if different granularity of reduction were desired for different roles.

This extension to role-symmetric games encompasses the broader class over which Ficici et al. [4] define the twins reduction. The clustering method by which they aggregate agents induces a role-symmetric game that restricts all roles to have the same strategy set (but allows different payoffs). They mention but do not develop the idea that the twins reduction might extend to role-symmetric games. To our

knowledge, hierarchical reduction has not been applied to role-symmetric games.

6. CONCLUSIONS

Our new player reduction method, deviation-preserving reduction, combines the most appealing aspects of hierarchical reduction and twins reduction. It also performs better than both prior methods experimentally: equilibria from *DPR* games have lower full-game regret and more closely resemble full-game equilibria, even when sampling fewer full-game profiles. In addition, performing iterated elimination of dominated strategies on deviation-preserving reduction games stays more faithful to the full game compared to other player reductions. Our alternative *DPR'* formulation performs reasonably well in the same tests. The simulation savings from *DPR'* are greatest when the reduced game has many players but few strategies, so *DPR'* may prove useful in such cases. Though it may not be obvious how to choose between *DPR* and *DPR'*, the evidence is quite compelling that deviation-preserving reduction is the best available player reduction method for analyzing large simulation-based games.

7. REFERENCES

- [1] B.-A. Cassell and M. P. Wellman. Agent-based analysis of asset pricing under ambiguous information. In *SpringSim Agent-Directed Simulation Symposium*, 2011.
- [2] P. Dandekar, A. Goel, M. P. Wellman, and B. Wiedenbeck. Strategic formation of credit networks. In *21st International Conference on World Wide Web*, Lyon, France, 2012.
- [3] A. Fabrikant, A. Luthra, E. N. Maneva, C. H. Papadimitriou, and S. Shenker. On a network creation game. In *22nd ACM Symposium on Principles of Distributed Computing*, pages 347–351, 2003.
- [4] S. G. Ficici, D. C. Parkes, and A. Pfeffer. Learning and solving many-player games through a cluster-based representation. In *24th Conference on Uncertainty in Artificial Intelligence*, pages 187–195, Helsinki, 2008.
- [5] P. R. Jordan, C. Kiekintveld, and M. P. Wellman. Empirical game-theoretic analysis of the TAC supply chain game. In *6th International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 1188–1195, Honolulu, 2007.
- [6] M. Kearns and Y. Mansour. Efficient Nash computation in large population games with bounded influence. In *18th Conference on Uncertainty in Artificial Intelligence*, pages 259–266, Edmonton, 2002.
- [7] K. Leyton-Brown and M. Tennenholtz. Local-effect games. In *18th International Joint Conference on Artificial Intelligence*, pages 772–780, Acapulco, 2003.
- [8] S. Phelps, M. Marcinkiewicz, S. Parsons, and P. McBurney. A novel method for automatic strategy acquisition in n -player non-zero-sum games. In *5th International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 705–712, Hakodate, 2006.
- [9] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
- [10] L. J. Schwartzman and M. P. Wellman. Stronger CDA strategies through empirical game-theoretic analysis and reinforcement learning. In *8th International Conference on Autonomous Agents and Multi-Agent Systems*, pages 249–256, Budapest, 2009.
- [11] G. Tesauro and J. L. Bredin. Strategic sequential bidding in auctions using dynamic programming. In *1st International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 591–598, Bologna, 2002.
- [12] G. Tesauro and R. Das. High-performance bidding agents for the continuous double auction. In *3rd ACM Conference on Electronic Commerce*, pages 206–209, Tampa, 2001.
- [13] P. Vytelingum, D. Cliff, and N. R. Jennings. Evolutionary stability of behavioural types in the continuous double auction. In *AAMAS-06 Joint Workshop on Trading Agent Design and Analysis and Agent Mediated Electronic Commerce*, Hakodate, 2006.
- [14] P. Vytelingum, D. Cliff, and N. R. Jennings. Strategic bidding in continuous double auctions. *Artificial Intelligence*, 172:1700–1729, 2008.
- [15] W. E. Walsh, R. Das, G. Tesauro, and J. O. Kephart. Analyzing complex strategic interactions in multi-agent systems. In *AAAI-02 Workshop on Game-Theoretic and Decision-Theoretic Agents*, Edmonton, 2002.
- [16] M. P. Wellman, D. M. Reeves, K. M. Lochner, S.-F. Cheng, and R. Suri. Approximate strategic reasoning through hierarchical reduction of large symmetric games. In *20th National Conference on Artificial Intelligence*, pages 502–508, Pittsburgh, 2005.