

Methods for Empirical Game-Theoretic Analysis (Extended Abstract)

Michael P. Wellman

University of Michigan
Computer Science & Engineering
Ann Arbor, MI 48109-2121 USA
wellman@umich.edu

Abstract

An emerging empirical methodology bridges the gap between game theory and simulation for practical strategic reasoning.

Game-Theoretic Analysis

Game-theoretic analysis typically takes at its starting point, most naturally, a description of its subject—the game, a formal model of a multiagent interaction. The recent surge in interest among AI researchers in game theory has led to numerous advances in game modeling (Gal & Pfeffer 2004; Kearns, Littman, & Singh 2001; Koller & Milch 2003; Leyton-Brown & Tennenholtz 2003) and solution techniques (Gilpin & Sandholm 2006; Porter, Nudelman, & Shoham 2004), substantially expanding the class of games amenable to computational analysis. Nevertheless, a great many games of interest lie well beyond the boundary of tractable modeling and reasoning. Complexity may be manifest in the number of agents or the size of their strategy sets (policy spaces), or the degree of incomplete and imperfect information. The issue here is not merely computational complexity of the analysis task (e.g., finding equilibrium), but actually the apparent impracticality of producing an explicit game model amenable to automated reasoning.

For example, consider the Trading Agent Competition Supply Chain Management (TAC/SCM) game (Eriksson, Finne, & Janson 2006). This is a well-defined six-player symmetric game of imperfect information, with interaction rules and exogenous stochastic processes described in a brief specification document. There is nothing particularly unusual about this game, nevertheless it presents a difficult challenge for game-theoretic analysis. The policy spaces and payoff functions are clearly induced by the specified rules, but the description is quite indirect. Even given complete policies for all six agents, there is no apparent means to derive the expected payoffs, short of sampling from the stochastic environment using an available game simulator. In this case, each sample point takes an hour to compute.

Copyright © 2006, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

Empirical Games

The approach we have been pursuing in my research group for the past few years¹ is to take the game simulator as the fundamental input, and perform strategic reasoning through interleaved simulation and game-theoretic analysis. The basic object of analysis is an empirical game, a description of the interaction scenario where payoff information comes in the form of data from observations or simulations. Constructing and reasoning about such games presents many interesting subproblems, which can be addressed by existing as well as new methods from simulation, statistics, search, and of course, standard game-theoretic analysis.

I find it useful to decompose empirical game-theoretic analysis into three basic steps. Many of the research contributions in this area manifest as techniques applicable to one of these subproblems, or results from approaches taken to them in a given domain.

Parametrize Strategy Space

Often the complexity of a game resides in vast policy spaces available to agents. Large spaces can arise, for example, from continuous or multi-dimensional action sets, as well as from imperfect information (when actions are conditioned on observation histories). It is often useful in such cases to approximate the game by restricting the strategy space, and structuring the space to admit a sensible search procedure. Results from analysis of restricted subgames often provide insight into the original game. Arguably, all applications of game theory in the social sciences employ stylized abstractions, which are manually designed restricted versions of actual games. From our perspective the interesting question is how to automate the abstraction process starting from a rich but intractable game specification.

One generic approach to generating candidate strategies is to start from some baseline or skeletal structure, and systematically introduce parametric variations. Some examples of natural baselines include:

1. Truthful revelation. For example, in an auction game, the baseline would be to bid one's true value. In the first-

¹Similar or identical techniques have also been employed by other researchers, especially those working experimentally with multiagent systems. Our main claim is in systematizing the methodology, in explicit game-theoretic terms.

price sealed-bid auction, this strategy guarantees zero surplus (!), but it turns out that the one-dimensional family of strategies defined by shading one's bid by a multiplicative factor includes excellent strategies, including the unique symmetric equilibrium (Reeves 2005).

2. Myopic best response. For example, in simultaneous auctions (SAAs), a natural starting point is straightforward bidding (Milgrom 2000), where the agent bids as though the current prices are final. We have explored an extensive family of bidding strategies for SAAs starting from this baseline, ultimately producing what we consider the leading contender in this domain (Osepayshvili *et al.* 2005).
3. Game tree search. The starting point for most programs designed to play complete-information turn-taking zero-sum games is minimax search. In a recent study of a 4-player chess game (Kiekintveld, Wellman, & Singh 2006), we defined the strategy space as a set of parametric variations on the basic game search architecture (e.g., control knobs for search depth and evaluation function weights).

Estimate Empirical Game

To illustrate some concepts associated with empirical games, we employ an example from a recent analysis of agents from the 2005 TAC/SCM tournament (Wellman *et al.* 2006). Figure 1 displays the empirical game, estimated from a sample of over 2000 game instances played with various combinations of six agent strategies. We describe the interpretation of this diagram in the course of explaining the game estimation and analysis.

Direct Estimation The most straightforward approach to estimate an empirical game from data is to treat the observations as direct evidence for the payoffs of the strategy profiles played. Toward this end we can bring to bear all the tools of Monte Carlo analysis, and related statistical techniques. We have found especially useful the method of control variates (L'Ecuyer 1994) for reducing variance based on adjusting for observable factors with known effects on payoffs. In the case of TAC/SCM, the most important factor is customer demand, which can significantly influence profits regardless of agent strategy. Applying control variates, we derive a measure of *demand-adjusted profit*, which we then employ as a proxy for payoffs in the empirical game estimation (Wellman *et al.* 2005a).

Each node in the graph of Figure 1 represents a profile of agent strategies. TAC/SCM is a 6-player symmetric game, and so with six possible strategies there are a total of $\binom{11}{5} = 462$ distinct strategy profiles to consider. We can reduce the game to a smaller version by requiring multiples of players to play the same strategy. Specifically, by restricting attention to cases where strategies are assigned to *pairs* of agents, we get an effective 3-player game, which we denote $\text{SCM}_{\downarrow 3}$. This game is combinatorially smaller, comprising only $\binom{8}{3} = 56$ profiles over the same 6-strategy set. The payoff to a strategy in an $\text{SCM}_{\downarrow 3}$ profile is defined as the *average* payoff to the two agents playing this strategy in the original 6-player game.

In several contexts, we have found experimentally and theoretically that this form of *hierarchical game reduction* produces results approximating well the original unreduced game, with great computational savings (Reeves 2005; Wellman *et al.* 2005b). Although we have not validated this specifically in TAC/SCM, intuitively we would expect that payoffs vary smoothly with the number of other agents playing a given strategy.

Our 2110 sample game instances (each requiring seven processor-hours to generate, not counting setup time and overhead due to failures) are distributed roughly evenly over the 56 $\text{SCM}_{\downarrow 3}$ profiles. In general, one may wish to deploy samples in a much more actively targeted manner. In other studies, we allocate samples with a view toward confirming or refuting candidate equilibria. The idea is to focus on promising profiles, and their *neighbors* in profile space—profiles that differ in the strategy choice of one agent. Walsh *et al.* (2003) have proposed information-theoretic criteria for selecting profiles to sample, and other approaches from Monte Carlo analysis and active learning should be applicable as well.

One special issue for empirical games is the need to handle partial coverage of observation data. Although in our illustrative example we have payoff estimates for all possible profiles, in many cases this will not be possible. We have found it useful in such cases to classify a profile s into one of four disjoint categories:

1. If the profile s has not been empirically evaluated (i.e., sampled), then we say it is *unevaluated*.
2. Otherwise, and for some neighbor t of s , the estimated payoff for the deviating agent is greater in t than s . In this case, we say s is *refuted*.
3. Otherwise, and some neighbor t of s is unevaluated. In this case, we say s is a *candidate*.
4. Otherwise, in which case we say s is *confirmed*.

In the empirical game of Figure 1, all profiles are evaluated. There is an arrow from each profile to its best deviation, which we define as the neighbor offering the greatest (positive) gain to the agent switching strategies to reach it. Deviations that are statistically significant are indicated by solid arrows. Since all nodes have outgoing arrows, we find that every profile is refuted in this empirical game, and therefore there are no pure Nash equilibria.

Regression An alternative approach to game estimation attempts to generalize the data across profiles, including those that are unevaluated in the sense of our classification above. The idea is to apply machine learning (regression) techniques to fit a payoff function over the entire profile space given the available data (Vorobeychik, Wellman, & Singh 2005). This approach offers the potential for reasoning over very large, even infinite, profile spaces.

Analyze Empirical Game

Analyzing an empirical game is much like analyzing any other; standard methods apply. Given the inherent uncertainty in reasoning about empirical games, it may be espe-

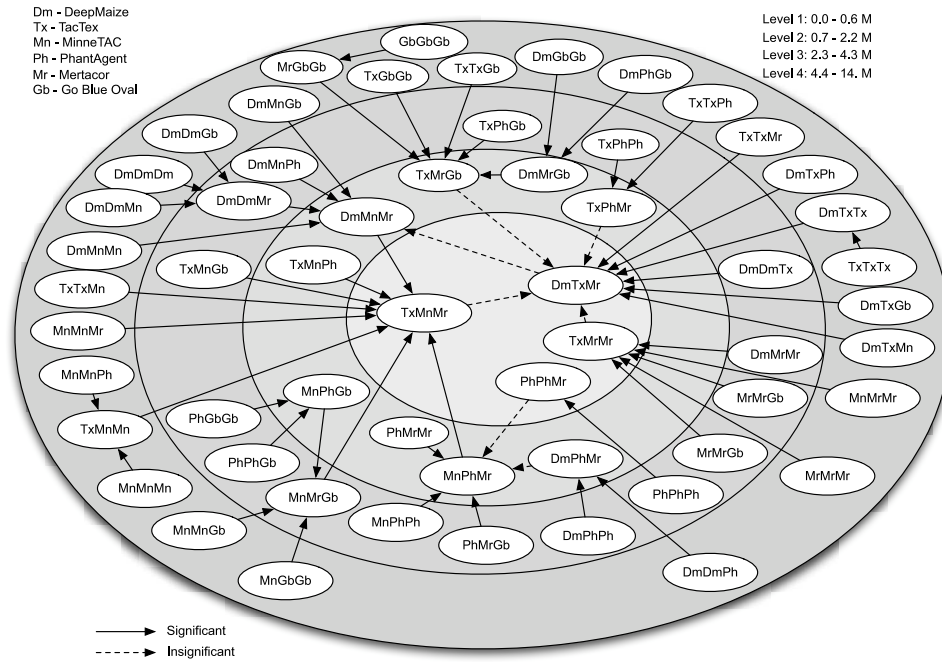


Figure 1: Deviation analysis of pure profiles of SCM_{3} .

cially appropriate to consider approximate equilibrium concepts, or more generally to reason about degrees of game-theoretic stability rather than categorical classes. Let $\epsilon(s)$ denote the maximum gain to deviating from s , over all agents and evaluated neighbors of s . If all neighbors are evaluated, then profile s is an $\epsilon(s)$ -Nash equilibrium, and for the special case $\epsilon(s) = 0$ it is a (confirmed) Nash equilibrium. If some neighbors are unevaluated, then $\epsilon(s)$ represents a *lower bound* on the ϵ rendering s an ϵ -Nash equilibrium.

The profiles in Figure 1 are arranged according to this ϵ measure. Profiles in the inner ellipse have $\epsilon \leq 0.6M$ ("M" represents a million dollars), with succeeding outer rings corresponding to increasing levels of this measure as indicated. With this interpretation, we can draw several conclusions by direct inspection of the diagram.

- Although no pure profiles are equilibria, some are much more stable than others.
- Each of four most stable profiles (and 11 out of the top 13) involve at least one agent playing the Mertacor strategy.
- Profiles where all agents play the same strategy (except PhantAgent) are among the least stable.
- Of the 35 profiles without Mertacor, 30 of them have a best deviation where some strategy changes to Mertacor. Of the 21 profiles with Mertacor, the best deviation changes from Mertacor in only three.

Of course, more specific analysis (e.g., deriving mixed-strategy equilibria) requires evaluation of the more precise payoff estimates.

Given finite data, it is also important to apply sensitivity analysis to the estimated game (Reeves 2005), or employ statistical bounds in reasoning about its implications (Vorobeychik, Kiekintveld, & Wellman 2006). The ϵ measure is useful for both purposes; we can derive distributions over $\epsilon(s)$ for each s , or provide confidence intervals with respect to $\epsilon(s)$. These measures can also provide guidance in sampling strategy, for example we might focus on refuting profiles with low $\epsilon(s)$. Finally, we have also employed the ϵ concept in evaluating techniques for payoff function regression (Vorobeychik, Wellman, & Singh 2005).

Applications

We have applied this methodology to a variety of games, especially market-based scenarios. In several cases we have been able to support conclusions about strategic issues in these games not accessible through standard analytic means.

- Verification that aggressive early procurement was a stable (but very destructive) behavior in the original (2003) TAC/SCM game (Arunachalam & Sadeh 2005), and that the preemptive policy of Deep Maize was an effective remedy (Wellman *et al.* 2005a).
- Identifying and validating a new strategy for dealing with the exposure problem in SAAs (Osepayshvili *et al.* 2005).
- Systematically evaluating strategies for the TAC travel game, leading to a choice that proved highly effective in TAC-05 (Wellman *et al.* 2005c).

Although much further work is required to develop the empirical game-theoretic approach into a precise, rigorous and fully automated methodology, the early results seem

quite encouraging. Further innovations in technique and experience in application promise to widen the scope of practical game-theoretic reasoning.

Acknowledgments

The work described here owes much to collaborations with Jeffrey MacKie-Mason and Satinder Singh, and our many student co-authors. This work was supported in part by the National Science Foundation under grants IIS-0205435 and IIS-0414710.

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.
- Gal, Y., and Pfeffer, A. 2004. Reasoning about rationality and beliefs. In *Third International Joint Conference on Autonomous Agents and Multi-Agent Systems*, 774–781.
- Gilpin, A., and Sandholm, T. 2006. Finding equilibria in large sequential games of imperfect information. In *Seventh ACM Conference on Electronic Commerce*.
- Kearns, M.; Littman, M. L.; and Singh, S. 2001. Graphical models for game theory. In *Seventeenth Conference on Uncertainty in Artificial Intelligence*, 253–260.
- Kiekintveld, C.; Wellman, M. P.; and Singh, S. 2006. Empirical game-theoretic analysis of chaturanga. In *AAMAS-06 Workshop on Game-Theoretic and Decision-Theoretic Agents*.
- Koller, D., and Milch, B. 2003. Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior* 45:181–221.
- L’Ecuyer, P. 1994. Efficiency improvement and variance reduction. In *Twenty-Sixth Winter Simulation Conference*, 122–132.
- Leyton-Brown, K., and Tennenholtz, M. 2003. Local-effect games. In *Eighteenth International Joint Conference on Artificial Intelligence*, 772–780.
- Milgrom, P. 2000. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy* 108:245–272.
- Osepayshvili, A.; Wellman, M. P.; Reeves, D. M.; and MacKie-Mason, J. K. 2005. Self-confirming price prediction for bidding in simultaneous ascending auctions. In *Twenty-First Conference on Uncertainty in Artificial Intelligence*, 441–449.
- 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.
- Reeves, D. M. 2005. *Generating Trading Agent Strategies: Analytic and Empirical Methods for Infinite and Large Games*. Ph.D. Dissertation, University of Michigan.
- 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*.
- Vorobeychik, Y.; Wellman, M. P.; and Singh, S. 2005. Learning payoff functions in infinite games. In *Nineteenth International Joint Conference on Artificial Intelligence*, 977–982.
- Walsh, W. E.; Parkes, D.; and Das, R. 2003. Choosing samples to compute heuristic-strategy Nash equilibrium. In *AAMAS-03 Workshop on Agent-Mediated Electronic Commerce*.
- Wellman, M. P.; Estelle, J.; Singh, S.; Vorobeychik, Y.; Kiekintveld, C.; and Soni, V. 2005a. Strategic interactions in a supply chain game. *Computational Intelligence* 21:1–26.
- Wellman, M. P.; Reeves, D. M.; Lochner, K. M.; Cheng, S.-F.; and Suri, R. 2005b. Approximate strategic reasoning through hierarchical reduction of large symmetric games. In *Twentieth National Conference on Artificial Intelligence*, 502–508.
- Wellman, M. P.; Reeves, D. M.; Lochner, K. M.; and Suri, R. 2005c. Searching for Wolverine 2005. In *IJCAI-05 Workshop on Trading Agent Design and Analysis*.
- Wellman, M. P.; Jordan, P. R.; Kiekintveld, C.; Miller, J.; and Reeves, D. M. 2006. Empirical game-theoretic analysis of the TAC market games. In *AAMAS-06 Workshop on Game-Theoretic and Decision-Theoretic Agents*.