RESEARCH STATEMENT

Patrick R. Jordan

March 27, 2009

My research addresses the problem of practical design and evaluation of complex scenarios involving multiple self-interested autonomous agents. The goal of my work is to develop computational methodologies and techniques that enable modelers to apply game-theoretic concepts to scenarios that defy classical analysis. In progressing towards this goal, my approach builds upon the fields of *computer science*, *economics*, and *operations research*.

In particular, I am concerned with algorithmic methods for analyzing market games, although the developed methods are largely applicable to general game-theoretic problems. Market games such as Internet ad auctions are pushing the boundaries of classical game-theoretic analysis and have spurred the development of new algorithmic analysis methods [9]. For scenarios that are too complex or lack directly specified game forms (i.e., payoff or utility functions), analytic game theory is not immediately applicable. In such situations, *empirical game models* [13], where observations or simulations of agent play are used to construct estimates of their utility, can support game-theoretic analysis despite lack of explicit game descriptions. Empirical game theory allows researchers to model a strategic scenario at a practical level of abstraction, where analysis is computationally feasible and game-theoretic concepts can still be applied. In the case of sponsored search, empirical game-theoretic modeling allows us to include many interesting aspects that have been traditionally excluded from analysis. In real-world implementations these aspects may play a pivotal role in advertiser and publisher strategies.

Empirical game-theoretic models are founded on an underlying game simulator, which generates outcomes from agent behaviors. The simulator defines the system in which the agents participate. Configuration parameters may include the number of agents, the allowable actions for each, possible type realizations (i.e., settings of private information for the agents), and other factors. Though the full strategy space allowed by the game simulator may be large or infinite, due to computational constraints empirical game models usually restrict the strategy space to a small number of heuristically defined strategies. Even within this restricted space, modeling accuracy is limited by statistical sampling, given the inherent stochastic behavior of typical simulation environments. Given a limited computational budget, modelers must choose carefully which strategy profiles to simulate and how many samples of those profiles to take.

One benefit of these empirical models is that game representations of a MAS can be made with varying degrees of fidelity to the underlying scenario. For instance, a stock market scenario may involve thousands of participants each of whom are selecting from a potentially infinite number of trading strategies. A simulation of the scenario may be reduced to a manageable number of representative participants each selecting from a handful of promising strategies. Depending on the computational budget, an analyst may increase or decrease the population size or available trading strategies in the simulator.

Empirical game-theoretic techniques have expanded researchers' capabilities of analyzing complex strategic scenarios by enabling increasingly faithful models of interaction. In my research, I focus on three general questions within empirical game theory. The questions arise from three different situations in which modelers may find themselves when approaching an empirical game-theoretic scenario:

1. Given an empirical model of the game, how should strategies be evaluated?
2. Given a set of heuristic strategies and a computational budget for simulating agent play, how should modelers optimally select profiles to be simulated?
3. Given observations of agent play, what game model best predicts payoffs resulting from agent play?

I elaborate on my approach to each of these research questions subsequently. In addressing these questions, I construct new methods for empirical analysis, which I use to address strategy and mechanism design concerns in complex market domains.

Evaluating Strategies in an Empirical Game Model. Standard game-theoretic analysis involves characterization of the *solutions* of the game, with the principal solution concept being the *Nash equilibrium*. More generally, we can analyze the *stability* of profiles of strategies in terms of *regret*. However, when we develop a similar notion of regret for strategies, we find that our evaluation is dependent on the choice of opponent strategies. My research in this area focuses on (a) developing a principled methodology for selecting opponent profiles and (b) pruning strategies through strategy elimination:

- **(a)** I introduce an equilibrium-based regret measure (*NE regret*) that evaluates a player's strategies in terms of its regret when faced with an opponent profile that supports a Nash equilibrium of the game. NE regret is used as the base evaluation criterion when ranking sets of strategies (*NE-response ranking*). Together with my coauthors, I analyze strategies [5] from the 2005 and 2006 TAC/SCM tournament, a market game of interest. Our analysis yields empirical evidence that agents are increasing in competency, in particular that they are improving responses to the previous year's equilibrium mixture. Using NE regret as a decision criterion, we explored the parameter space of our agent in the TAC/SCM competition. This led to a first-place finish in the 2008 TAC/SCM tournament as well as general insights regarding different elements of agent strategy in the TAC/SCM game.

- **(b)** Eliminating strategies can be used as a pre-processing step to reduce the complexity of many game-theoretic algorithms, such as finding equilibria, a central component of NE-regret analysis. General eliminability criteria, like that introduced by Conitzer and Sandholm [3], can greatly reduce the size of strategy space. However, this comes at a price. While general elimination criterion may not rule out equilibrium strategies in the original game, it may introduce NE in the reduced game that are not equilibria in the original game. This is undesirable, we would like to know that profiles identified as NE in the reduced game persist as NE in the original game. Persistence of this type is very important for NE-regret computations and, in general, sets of strategies that are *closed under rational behavior* (CURB) [1] retain this property. Benisch et al. [2] introduce a set of algorithms for finding minimal CURB sets in two-player games and analyze their complexity, unfortunately for three or more players the feasibility test that is central to the algorithm no longer remains linear. I have developed an extension of their algorithm for general n-player games that identifies the *primitive formation* sets of Harsanyi and Selten [4], which preserves the linearity of the feasibility test. The primitive formation sets weakly contain the minimal CURB sets.

When faced with the high computational cost of simulation, techniques that reduce the size of the game space become particularly valuable. One may be willing to forgo exact solutions, for instance exact Nash equilibria, for approximate solutions like ϵ-Nash equilibria, if the size of the game can be greatly reduced for analysis. I am currently pursuing efficient techniques for determining minimal ϵ-formations in games, given constraints on the size of the reduced game. Another interesting extension of my work regards the analysis of strategies in games where the payoffs are partially specified, as is often the case with empirical game-theoretic analysis.

Selecting Profiles for Simulation. The ability to apply game-theoretic analysis techniques to a scenario relies upon the existence of an empirical game model, which in turn relies on a set of observations used to create that model. I have developed and analyzed approaches to selecting a set of profiles to be observed when we are constrained by some computational budget. The core of the approach transforms the *profile selection problem* into a *profile search problem* with each action in the search corresponding to an observation of a profile. The given search algorithms are designed with the goal of identifying approximate equilibria, which is motivated by the pivotal role these profiles play in game-theoretic analysis and mechanism design. The justification for framing the selection problem as a search problems forms around two key observations. First, simulation is costly, often the dominant cost in analysis of empirical games. Second, observing the entire game 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. If an (approximate) equilibrium can be identified without simulating the entire space, this can result in a substantial computational savings. I have developed search algorithms for efficiently searching for approximate equilibria in empirical games [6] under two distinct models of observation.

- **Revealed-payoff model:** In this model, each observation determines the true payoff for a designated pure-strategy profile. Sureka and Wurman [10] propose an algorithm based for this model based on tabu best-response search (TABU) to search for pure-strategy Nash equilibria within the profile space. I compare TABU to minimum-regret-first search (MRFS), employed by Vorobeychik at al. [11], and find that MRFS significantly outperforms TABU in terms of its ability to confirm better approximate equilibria with fewer profile simulations.

- **Noisy-payoff model:** In this model, an observation is a stochastic sample of the payoff for a designated pure-strategy profile. I introduce a new algorithm based on information gain, called IGS, and find that it outperforms the ECVI repeated sampling algorithm, the current benchmark in the literature [12]. Unlike ECVI, which was an attempt to construct a computationally feasible version of expected value of information, the IGS family of algorithms does not directly resolve to improve the mean estimate of player regret. Instead, the IGS algorithm focuses on improving some distribution over profiles. For example, this distribution could be a distribution of play, a likelihood of pure-strategy Nash equilibrium, or the probability that a profile minimizes regret. Optimizing the distribution rather than a point estimate can improve calculations involving the distribution and other heretofore unknown quantities. Consequently, an expected-value-of-information-based algorithm like ECVI may not completely capture the decision-theoretic problem underlying the game analysis task.

So far, I have focused on developing search algorithms that use a fixed underlying game form as a basis for selecting profiles. If, for instance, a compact representation for a game was known, this could significantly reduce the amount of simulation require to accurately estimate the payoffs of the game. Therefore, one important future area of research considers methods for payoff selection that efficiently learn the structure of a game.

Choosing the Best Game-Theoretic Model for a Scenario. In much of the previous research on empirical games, the strategy space of the empirical game is assumed identical to that of underlying simulation. We can relax this constraint, allowing models where the strategy sets or even players do not correspond precisely to the base notions defined by the simulator. For instance, an empirical game model may treat two strategies that are distinct for the simulator as interchangeable in its own strategy space. This coarsens the model the empirical game uses to predict payoffs, reducing the model's complexity compared to the finer-grained strategy space of the simulator.

Entertaining multiple candidate models of varying complexity provides useful flexibility. A more complex model may capture observations better than a simpler one, however it may also be more susceptible to fitting spurious information in the observations. I term this the *generalization risk* associated with an empirical game model. In the game-theoretic context, the consequence of incorrect generalization could be that profiles that appear stable are actually unstable, or vice versa. In recent work [8], I describe a measure for generalization risk in empirical game and algorithms for selecting an optimal model. The generalization risk framework provides a general criterion for empirical modeling choices, such as adoption of factored strategies or other structured representations of a game model.

Practical Strategic Analysis of Market Games. Empirical game-theoretic analysis (EGTA) has proven to be a valuable tool in reasoning about strategies and mechanisms in complex market scenarios that defy classical analysis. As new markets and technologies emerge, algorithmic solutions that can enable practitioners to make principled trade-offs between model fidelity and computational feasibility will take a central design role. I highlight two market games of interest in which empirical game-theory has and can make significant contributions:

- **Supply Chain Management:** Developing a strategy to manage a supply chain is a complex task. Top level features of a strategy often include decision coordination mechanisms, choice of optimization algorithms, and modeling choices for dealing with uncertainty. These features are typically parameterized. Setting these parameters is a decision process that should be reasoned about strategically. EGTA has shown to be an efficacious methodological process for parameter selection that can scale with computational resources.
- **Ad Auctions:** The emergence of Internet advertising, specifically ad auctions, as a significant commercial success over the past decade has led to increasing interest among academic researchers, manifest in a growing literature and a popular regular workshop on the topic. Analytical approaches have provided a strong theoretical foundation for understanding the strategic nature of the auctions, however potentially important aspects of the auctions are often abstracted in order to make the model feasible for analysis. My co-developers and I propose a new research competition to analyze some of these overlooked aspects [7]. Both the commercial importance and academic interest were major motivations of introducing a new Trading Agent Competition (TAC) game in this area. Given that bidding in keyword auctions (employing essentially the same mechanism we incorporate in the game) is a widespread current activity, the prospects for real-world implementation of ideas developed in the research competition are more direct than previous TAC games.

Summary My goal as a researcher is twofold. I attempt to expand the set of tools and techniques available for algorithmic and empirical game-theoretic analysis. In particular, I endeavor to find efficient methods for constructing and analyzing games with respect to an underlying simulator. In a complementary pursuit, I seek to design optimal strategies and mechanisms for real-world economic applications.

References

[1] Kaushik Basu and Jörgen W. Weibull. Strategy subsets closed under rational behavior. *Economic Letters*, 36(2):141–146, 1991.

[2] Michael Benisch, George B. Davis, and Tuomas Sandholm. Algorithms for rationalizability and CURB sets. In *Twenty-First National Conference on Artificial Intelligence (AAAI)*, Boston, July 2006.

[3] Vincent Conitzer and Tuomas Sandholm. A generalized strategy eliminability criterion and computational methods for applying it. In *National Conference on Artificial Intelligence (AAAI)*, pages 483–488, 2005.

[4] John C. Harsanyi and Reinhard Selten. *A General Theory of Equilibrium Selection in Games*. MIT Press, June 1988.

[5] Patrick R. Jordan, Christopher Kiekintveld, and Michael P. Wellman. Empirical game-theoretic analysis of the TAC supply chain game. *Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 1188–1195, May 2007.

[6] Patrick R. Jordan, Yevgeniy Vorobeychik, and Michael P. Wellman. Search for approximate equilibria in empirical games. *Seventh International Joint Conference on Autonomous Agents and Multi-Agent Systems*, 2:1063–1070, May 2008.

[7] Patrick R. Jordan and Michael P. Wellman. The Trading Agent Competition Ad Auctions game. `http://aa.tradingagents.org`, 2009.

[8] Patrick R. Jordan and Michael P. Wellman. Generalization risk minimization in empirical game models. *Eighth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, To Appear.

[9] Sébastien Lahaie, David M. Pennock, Amin Saberi, and Rakesh V. Vohra. Sponsored search auctions. In *Algorithmic Game Theory*, pages 699–718. Cambridge University Press, 2007.

[10] Ashish Sureka and Peter R. Wurman. 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*, pages 1023–1029, New York, NY, USA, 2005. ACM Press.

[11] Yevgeniy Vorobeychik, Christopher Kiekintveld, and Michael P Wellman. Empirical mechanism design: Methods, with application to a supply chain scenario. *Seventh ACM Conference on Electronic Commerce*, pages 306–315, June 2006.

[12] William Walsh, David Parkes, and Rajarshi Das. Choosing samples to compute heuristic-strategy Nash equilibrium. In *Fifth Workshop on Agent-Mediated Electronic Commerce*, 2003.

[13] Michael P Wellman. Methods for empirical game-theoretic analysis. In *Twenty-First National Conference on Artificial Intelligence (AAAI)*, pages 1152–1155, Boston, July 2006.