Œconomia
History, Methodology, Philosophy

10-2 | 2020
Simulations in Economics: Methodological and Historical Perspectives

Economic Reasoning from Simulation-Based Game Models

Le raisonnement économique dans les modèles de jeux basés sur des simulations

Michael P. Wellman

Electronic version
URL: http://journals.openedition.org/oeconomia/8386
DOI: 10.4000/oeconomia.8386
ISSN: 2269-8450

Publisher
Association Œconomia

Printed version
Date of publication: 1 June 2020
Number of pages: 257-278
ISSN: 2113-5207

Electronic reference
Michael P. Wellman, « Economic Reasoning from Simulation-Based Game Models », *Œconomia* [Online], 10-2 | 2020, Online since 01 June 2020, connection on 06 August 2020. URL : http://journals.openedition.org/oeconomia/8386 ; DOI : https://doi.org/10.4000/oeconomia.8386

Economic Reasoning from Simulation-Based Game Models

Michael P. Wellman*

Simulation modeling in economics has historically been viewed as an alternative to mainstream analytic technique, and as such has generally and intentionally avoided the focus on rational behavior and equilibrium reasoning that is characteristic of game-theoretic approaches. The emerging methodology of empirical game-theoretic analysis (EGTA) attempts to bridge agent-based simulation and game theory, combining the flexibility of simulation with the discipline of rationality criteria. The basic idea is to estimate a restricted game model from data generated by systematic simulation over a space of heuristic strategy profiles. EGTA enables a standard form of strategic analysis for complex economic scenarios previously addressed only in severely stylized form.

Keywords: agent-based modeling, agent-based simulation, empirical game-theoretic analysis, computational game theory

Le raisonnement économique dans les modèles de jeux basés sur des simulations
Historiquement, la modélisation par simulation en économie s'est développée comme alternative aux techniques analytiques courantes. De ce fait, elle a généralement et sciemment évité de traiter de la rationalité et du raisonnement à l'équilibre, qui sont au cœur de la théorie des jeux. La nouvelle méthodologie d'analyse empirique de jeux (EGTA en anglais) tente de faire le pont entre la simulation à agents multiples et la théorie des jeux, combinant la flexibilité de la simulation avec la discipline des critères de rationalité. L'idée principale est d'estimer un modèle de jeu réduit à partir de données générées par des simulations systématiques sur un espace heuristique des profils de stratégie. L'EGTA fournit une norme stratégique pour analyser les scénarios économiques

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

I am grateful for comments from members of the Strategic Reasoning Group at the University of Michigan, participants in the *Œconomia* workshop held in Paris in September 2019, and two anonymous referees.

complexes qui étaient auparavant modélisés uniquement sous des formes fortement stylisées.

Mots-clés: modélisation multi-agents, simulation multi-agents, théorie des jeux, économie computationnelle
JEL: C63, C7, D44

Advocates of agent-based modeling (ABM) in economics (Fagiolo and Roventini, 2012; Gilbert, 2007; Miller and Page, 2007; Richiardi, 2012; Tesfatsion, 2006) tend to invoke two kinds of arguments in favor of simulation-based approaches relative to prevailing analytical methods.

1. Simulation can accommodate a great deal of complexity, in comparison to models amenable to theoretical characterization. For example, heterogeneity in agent types, partial and asymmetric information, and dynamic interaction all pose significant challenges for analytic reasoning, but can be incorporated in simulation with little additional effort. In practice, an imperative for theoretical characterization therefore limits model complexity, potentially in ways that systematically distort the kinds of issues addressed or phenomena recognized.

2. Standard solution concepts impose unrealistic rationality requirements or adhere too strictly to equilibrium assumptions. A "bottom-up" (Epstein and Axtell, 1996) approach to generating economic outcomes may be more likely to produce novel insight (particularly on "emergent" phenomena) than models conceived in a "top-down" manner. As Richiardi (2018) points out, whereas equilibrium reasoning can be relevant to bottom-up systems, it is problematic to ignore the dynamics of adaptation or other out-of-equilibrium behavior.

Though these arguments are typically put forth together and viewed as complementary, let us take them separately for current purposes. The first is a purely technical point and perhaps uncontroversial. Whereas clever theorists continually extend the boundary of analytic tractability, it is hard to deny that complexity is a constraint, and that economically relevant features are often sacrificed for the sake of maintaining feasibility of mathematical treatment. It is far easier to include complicating features in a simulation model. Though complexity may later impose a burden in understanding simulation results, it generally does not add significant cost to development and execution of the simulation model itself.

The second category of argument is less straightforward. Concepts like rationality and equilibrium may be restrictive, but that restrictiveness serves to impose a necessary discipline on agent behaviors considered. Whether these concepts bring exactly the right form of restrictiveness is debatable, but the need for *some* discipline is compelling, as virtually any outcome may be producible by some pattern of agent behavior. Indeed, the degree of freedom afforded the designer of agent-based simulations is perhaps the greatest cause for skepticism about simulation results. The way that some (though by no means all) ABM advocates have positioned their methods in opposition to mainstream economics has probably also contributed to the lack of general acceptance of these techniques in the field to date.

For the past 20 years or so, some computer scientists have been pursuing an approach that aims to exploit the advantages of simulation for modeling complex systems while maintaining the appeal to selection of rational agent behavior. The approach is a hybrid of ABM and game theory called *empirical game-theoretic analysis* (EGTA) (Tuyls et al., 2020; Wellman, 2006; 2016). In the sequel, we explain the key components of EGTA methodology and demonstrate through examples its value for economic reasoning. We also discuss the challenges and limitations of this approach, and its technical and historical relations to ABM and other simulation approaches in economics.

1 EGTA: Basic Concepts

The defining feature of EGTA is induction of a game model from systematically generated simulation data. Each run of the simulator produces data about the outcome of the game, including payoffs for each of the players. The simulator can itself be viewed as a model of the game, albeit one expressed in procedural form not directly amenable to game-theoretic analysis. Typically we start from a set of candidate agent behaviors, that is, *strategies* for playing the game modeled by the simulator. These strategies may represent heuristic policies for the game, based on behavioral data, approximate optimization, or any other approach typically employed in agent-based modeling. Often the strategies are points in a well-defined space of feasible policies, specified in terms of tunable parameters or other controllable features.

1.1 A Very Simple Example

Let us illustrate with a toy example. Consider a game between two bidders in a first-price sealed bid (FPSB) auction.[1] We assume independent private values, distributed uniformly on [0,1]. For candidate strategies, we consider those where the players bid a constant fraction k of their private value, with $k \in \{1/3, 1/2, 2/3\}$. Thus we have a 3×3 two-player game. Simulating this game is trivial: All we do is draw private values for the two players, apply their bidding strategy, and declare the highest bidder the winner. If a player

1 This game is simple enough to characterize analytically (Krishna, 2010), so one would certainly not employ EGTA in this instance. Nevertheless, examining a game with known solution helps to illustrate the concepts, and as even small variants on this game may be analytically intractable, it well exemplifies some contexts where EGTA has been applied.

with strategic parameter k and value v wins, it obtains profit $(1-k)v$. The losing bidder gets zero profit.

Results from the exercise of estimating the game in this way are shown in Table 1. At left is the payoff table for the exact 3×3 normal-form game, taken from analytically derived formulas (Schvartzman and Wellman, 2009a). At right is an empirically derived model of the game, estimated by sampling. With 500 samples per strategy pair (for symmetric pairs we get two payoff observations per sample, for a total of 1000), the estimates exhibit errors in the second or third decimal place. For a game this simple, it would be feasible to take millions of samples and get quite a few more decimal places right. When simulation of a game instance is more computationally intensive, however, or when there are many more strategy combinations to be considered, sampling errors of this sort may be expected in empirical game estimation.

Table 1: Normal Forms for Two-Player FPSB with Three Heuristic Strategies.

k	1/3	1/2	2/3
1/3	2/9	4/27	1/9
1/2	11/54	1/6	1/8
2/3	11/72	13/96	1/9

k	1/3	1/2	2/3
1/3	.2234	.1252	.1198
1/2	.2258	.1659	.1210
2/3	.1405	.1371	.1066

Payoffs are shown for the row player. Left: exact payoffs calculated analytically. Right: empirical game estimated by numeric simulation, 500 samples of each strategy pair.

Even with such errors, in this case the unique solution of the empirical game, $(1/2, 1/2)$, is also the unique Nash equilibrium of the *underlying* (true) game. Not all game-theoretic inferences correctly transfer, however: For instance, $k = 1/2$ is a dominant strategy in the empirical game but not in the underlying game.

1.2 Parametrized Heuristic Strategies

A key design decision in any EGTA study is what strategies to make available to the respective players. Inevitably, the strategy sets will be strict subsets of the full set of behaviors possible in the underlying game. In the spirit of agent-based modeling, strategies tend to be *heuristic*, embodying insights about effective behavior drawn from intuition or experience. Heuristic strategies may even incorporate optimization techniques or represent optimal behavior in idealized conditions, but generally fall short of exact equilibrium solutions (if we could solve for equilibria we would not need to adopt a simulation-based approach). In some cases, modelers may prefer to limit the considered strategies to reflect bounded rationality considerations or observed characteristics of real-world behaviors.

Since heuristic strategies are by nature imprecise fits for particular circumstances, common practice is to expose control parameters by which the heuristics may be tuned to specific environments. In our FPSB example, there is a single strategic parameter k, representing the fraction of valuation the agent will bid. In this instance, the parameter space happens to cover the exact equilibrium, though for typical subject games this would not be possible to ensure in practice. The k parameter in any case captures an intuitive way to explore a natural range of behaviors, from lowball bidding ($k \to 0$) to conservatively offering near one's valuation ($k \to 1$). The setting $k = 1$ also corresponds to the concept of truthful bidding, which is a poor strategy in FPSB (guaranteed zero surplus) but represents a natural landmark for bidding problems in general.

For auctions and other games related to revelation mechanisms, degree of deviation from truthfulness (i.e., *bid shading*) is a natural candidate for a strategy parameter. In our FPSB example, the quantity $1 - k$ represents the shading proportion, and more broadly, shading parameters are typical in EGTA studies of competitive bidding. For example, Zhan and Friedman (2007) performed a simulation-based game-theoretic analysis of double-auction markets, with buyer strategies defined by shading parameters, and seller strategies inversely by *markup* parameters. This study considered additive and exponential in addition to multiplicative application of shading (and markup), further refined in follow-up work by Cervone et al. (2009).

More complex strategies will naturally admit more elaborate parametrizations, ranging over multiple dimensions. For example, Cliff (2006) conducted extensive evolutionary optimization of double-auction bidding strategies over 60 parameters. These parameters controlled aspects of direct behavior such as tendency to shade bids, as well as higher-order factors like rate of adaptation (learning). For environments with multiple forms of decisions (e.g., interactions across markets with distinct mechanisms), strategies may employ distinct approaches for each form, leading to multifaceted parameter spaces (Wellman et al., 2006). Once behavior is complex enough, agent designers find it useful to describe the fixed or structural elements of this behavior as the agent *architecture* (Weiss et al., 2010), naturally leaving the variable elements to be exposed as control parameters.

1.3 Incomplete Game Models

A natural measure of the complexity of a normal-form game is the number of strategy combinations, or *profiles*, that is, ways of mapping the players to chosen strategies. For a 3×3 game like the toy FPSB example above, it is quite feasible to evaluate through simulation all strategy profiles. More generally, the size of the profile space is exponential in the numbers of agents and strategies, and so exhaustive estimation will typically be infeasible. Fortunately, in many cases, game-theoretic solutions can be characterized far short of simulating the entire game. For example, after analyzing the empirical FPSB game of Table 1, we may wish to test the provisional solution by exploring additional strategies, say k' and k''. Rather than simulate these strategies against all the previous strategies (and each other), we might just test them as deviations from the $(1/2, 1/2)$ equilibrium. As long as neither $k = k'$ nor $k = k''$ outperform $k = 1/2$ against the other bidder playing 1/2, we may conclude (modulo sampling error) that the original solution remains a solution to the extended game. This leaves open the possibility that other solutions exist, and indeed we cannot reach conclusions about the full set of solutions without exhaustive evaluation, unless we invoke additional game-specific assumptions.

We refer to a game model as *incomplete* if only a strict subset of the strategy profiles have been evaluated through simulation.[2] Researchers have developed a variety of approaches to explore and characterize incomplete game models. This includes methods to drive simulation toward regions of the game model that are deemed most important. Sureka and Wurman (2005) proposed a procedure for extending an incomplete game model based on tabu best-response search. Subsequent works identified and evaluated alternative search procedures, likewise under a model where the basic operation is evaluating the expected payoff of a specified pure profile (Jordan et al., 2008; Vorobeychik et al., 2006). Fearnley et al. (2015) formalized this *payoff query* approach to game solving and characterized situations where approximate equilibria may be identified in the worst case without exhaustive exploration of the profile space.

Perhaps the simplest illustration of profile-space exploration is a procedure that Wellman et al. (2013) refer to as the "EGTA inner loop" (see Figure 1). Let us define a *candidate* solution to be a strategy profile (generally, mixed) that is evaluated in the current incomplete game, but such that there is no strategy for any player that is a beneficial deviation in the incomplete game. Since the game is incomplete, a potential deviation may be unevaluated (i.e., some profiles required to calculate the deviation payoff have not been simulated), or if evaluated, may or may not be beneficial. If all potential deviations have been evaluated and are non-beneficial, then the candidate is *confirmed* as a Nash equilibrium. Otherwise, if some deviation is unevaluated, the candidate is *unconfirmed*. The EGTA inner loop, in somewhat simplified form, repeats the following steps until **done**:

1. Analyze the incomplete game to identify candidates.

2. If there exists an unconfirmed candidate, then simulate all unevaluated deviations, and go to step 1.

3. If there exists a confirmed candidate, then **done**. Otherwise, find a subgame with all profiles evaluated, extend it by adding a strategy and completing the extended subgame, and go to step 1.

2 Note that even a "complete" game model by this definition is a restricted version of the original game of interest, since the strategies considered are only enumerated sets of parametric instances of the full space of possible strategies.

Note that even when the inner loop is "**done**" and we have a confirmed solution candidate, there remains opportunity to refine the model by considering additional strategies. This is reflected in the "outer loop" of Figure 1, in which additional strategies are brought into consideration based on results from analyzing the current empirical game model.

Figure 1: Iterative Procedure for Reasoning about Incomplete Empirical Game Models.

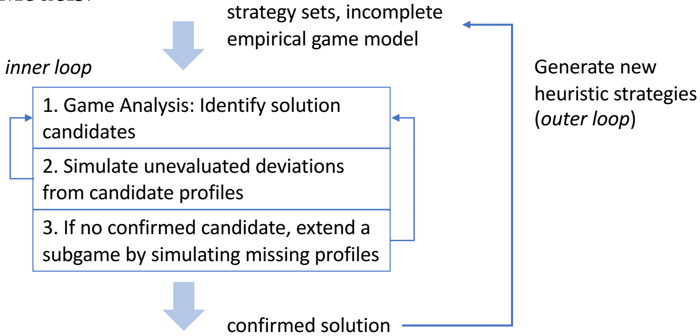

As illustrated in the presentation of the toy FPSB example above, simulation does not exactly reveal expected payoffs, but rather noisy samples as a means to estimate such payoffs. Hence the computational resources available for sampling need be allocated considering not only which profiles to evaluate, but according to the relative accuracy needed among those profiles evaluated. This issue was addressed directly in early EGTA work by Walsh et al. (2003), who introduced a criterion they termed *expected confirmational value of information*. Alternative criteria were later proposed by Reeves et al. (2005) and Jordan et al. (2008). Subsequent work has also addressed statistical characterization of EGTA results (Vorobeychik, 2010), including direction of sampling to maximize statistically valid conclusions (Jecmen et al., 2020; Wiedenbeck et al., 2014).

Another approach is to go beyond estimation of profile payoffs by direct simulation of those profiles, and instead fit a general game model to whatever simulation data has been collected. Given payoff data generated over a range of sampled profiles, induction of a game model can be viewed as a statistical machine learning problem. A variety of methods for learning game models from data have been developed (Honorio and Ortiz, 2015; Li and Wellman, 2020; Sokota et al., 2019; Vorobeychik et al., 2007; Wiedenbeck et al., 2018). Basing a game model on data—even simulated data—may overcome skepticism that some harbor about game-theoretic models.

2 EGTA and Agent-Based Modeling

Strictly speaking one should classify EGTA as a special case of ABM, in that EGTA methodology entails simulation of strategic behavior, effectively employing a model which would naturally be labeled as “agent-based”. This label does not have an unambiguous technical definition, but generally is understood to focus on entities that autonomously generate behavior and can be ascribed *agent attitudes* such as beliefs, goals, and intentions. There is no requirement that strategies in an agent-based simulation explicitly reference such attitudes, and practice within the ABM genre varies in extent of appeal to explicit beliefs and preferences.[3] As Schinckus (2019) notes, there are diverse methodological perspectives motivating agent-based models, and these may be more or less complementary with (versus substitutable for) standard agent concepts. In particular, the ABM literature includes an extensive body of simulation studies (roughly aligned with so-called “EconoPhysics” perspectives) where the actors have simple particle-like behaviors and interactions. More complex strategic domains would tend to call for more sophisticated agent designs, better aligning with senses of the term “agent” from the field of artificial intelligence (Wellman, 2016).

The focus on simulation as the tool for examining agent interactions establishes a foundation of shared elements between ABM and EGTA methodology. Figure 2 shows how the methodologies relate, including these common features as well as more distinguishing characteristics. As discussed in the introduction, the foremost motivation for simulation is its ability to accommodate *complexity* in the environment. Complexity may manifest, for example, in nonlinear effects of actions, interactions dependent on fine-grained or localized features, path-dependent dynamic patterns, and asymmetric information structures. Environments with such features rarely admit closed-form or provably optimal strategies, hence the necessity of *heuristic* behavior specification. The appeal of heuristics also lies in the ability to incorporate in procedural form ideas about how agents make decisions, based on empirical evidence or behavioral theories.

The agent-based modeling approach also fosters consideration of *heterogeneity* among agents. Members of a population invariably differ in beliefs (e.g., based on differences in available information) and preferences, as well as capabilities and opportunities, among other things. These differences may be parametric or categorical, and arise randomly or through a generative process (or some combination of the above). Regardless, the heterogeneity itself is often pivotal in driving the trajectory of the multi-agent system, and so approaches based on aggregate or representative-agent behavior may be fundamentally inadequate (Chapman and Polkovnichenko, 2009).

Where EGTA departs from the main line of ABM research is in how it explores and selects among candidate agent designs from which to draw its conclusions. EGTA is motivated by the goal of understanding rational behavior in a complex strategic environment modeled by a simulator. The “GT” in EGTA stands for “game-theoretic”, after all, which naturally means

3 Though since the target of simulation in EGTA is to build a game model there must at least be some explicit quantification of outcome *payoffs*. Such payoffs technically correspond to agent utility functions, representing the agents’ preferences over the outcome space.

Figure 2: Comparison of ABM and EGTA Methods and Practice.

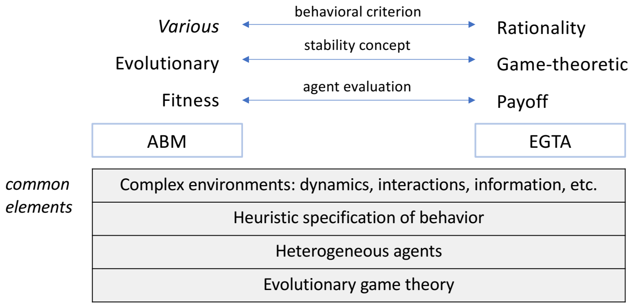

The foundation of common elements derives from the methodologies' central use of simulation. The crux of differences lies in the criteria for selecting among agent behaviors. ABM employs various criteria, often informal, whereas EGTA expressly filters for rational configurations based on game-theoretic solution concepts.

that it will appeal to GT stability concepts (in turn, derived from rationality theories) in assessing the relevant agent strategies.

ABM studies in economics also sometimes apply stability concepts, though these tend to be based on evolutionary dynamics rather than rationality criteria. This stems in part from the way the agent-based approach has been motivated as a departure from mainstream neoclassical frameworks, and especially in contrast to the prevailing focus on equilibrium reasoning. Evolutionary thinking in economics, as pioneered by Nelson and Winter (1982), emphasizes change in behavior through incremental adjustment and feedback over extended and repeated interactions. This is framed in explicit contrast to optimization-oriented models, such as commonly invoked in mainstream economic theory and inherent to game-theoretic treatments.

However, avoiding explicit optimization does not obviate the need for some way to evaluate agent performance. Following the influence of theoretical work in evolutionary biology, evolutionary modeling in ABM measures success of an agent by *fitness*, a term meant to capture tendency to thrive or grow in prevalence within a population of evolving entities. In ABM, fitness typically represents the same sort of outcome features (e.g., profit in an economic context) that would be captured by a payoff or utility function. As such, the difference is more terminological than substantive. Based on the chosen evaluation function, ABM studies may include an evolutionary tournament whereby a population of agent behaviors is iteratively adjusted according to relative fitness, for example using the model of *replicator dynamics* (Schuster and Sigmund, 1983). The evolutionary path in this approach is computed by simulation of a population over a series of generations, ultimately converging to a set of strategies that is *evolutionarily stable* (Taylor and Jonker, 1978). For large populations or sets of candidate behaviors, the associated computational burden may limit consideration to fairly simple agent behaviors and fitness functions.

Concern with evolutionary dynamics and outcomes is also shared by a significant thread of work in EGTA (Phelps et al., 2005; Tuyls et al., 2020). Tuyls and Parsons (2007) argue that *evolutionary game theory* is particularly relevant to understand systems of learning agents, and that analyzing the trajectory of agent behavior under evolutionary dynamics can provide special insights. Studies that perform evolutionary analysis under the EGTA label contrast algorithmically with typical ABM practice by separating the simulation-based generation of fitness data from the simulation of evolution itself. EGTA practice is to estimate the fitness of a space of strategy profiles through simulation, constructing a *heuristic payoff table* (essentially estimating the empirical game), which in turn provides sufficient data for evolutionary computation without requiring further simulation of the agents themselves (Bloembergen et al., 2015).

From a broader perspective, evolution and game theory are two particular approaches to characterize stability of a configuration of agent behaviors. As such, they provide alternative criteria for selecting among candidate configurations. In principle, one could apply any sufficiently well-defined criterion in a systematic manner for behavior selection, just as evolutionary and rationality-based solution concepts are employed in EGTA. ABM practice often promotes alternative criteria, such as generating behavior that reproduces stylized facts of a domain. Such stylized facts are not always sufficiently formalized to automate a systematic search, but nevertheless do provide an effective way to filter candidate behaviors.

3 Historical Development

The first paper explicitly describing an EGTA methodology was written by a group of IBM researchers (Walsh et al., 2002), advocating equilibrium (strategic and evolutionary) analysis of games estimated by simulation over heuristic strategies. This early work studied a dynamic pricing game with 5–20 agents over three heuristic strategies, as well as a double-auction game over 14 or 20 agents, again with three heuristic strategies. As described in the paper, the main idea is to compute through simulation what the authors term a *heuristic-payoff table* representing expected payoffs over the heuristic profile space.

The approach of Walsh et al. (2002) was an outgrowth of prior agent-based modeling efforts by themselves and others in these domains, in particular on the dynamics of interacting heuristics for automated shopping and pricing (Greenwald et al., 1999). These works in turn extended an emerging practice in agent-based modeling to explore various combinations of behavior, and often to trace out the dynamics of evolutionary or otherwise adaptive variation. For example, adaptive variation was a key feature of early agent-based studies of financial markets (Arthur et al., 1997).

Moving from evolutionary dynamics and stability concepts to game-theoretic stability is in some respects a small conceptual step; however, it is not one that the field of agent-based economics would have been inclined to take. As discussed above, agent-based modeling was seen by many practitioners as an alternative to analysis based on rational behavior, so applying game theory to data from agent simulations might have been considered anathema. The work by Greenwald and Kephart (1999) was different in that it expressly appealed to game-theoretic concepts in the definition of heuristic strategies and characterization of steady-state behavior. However, the step of applying game-theoretic analysis to the interactions among heuristic strategies was taken for the first time in the paper by Walsh et al. (2002).[4]

Trading and bidding games, particularly double-auctions, were popular targets for agent-based simulation game modeling in the earliest days of its development. Phelps et al. (2005) employed the approach to compare two alternative double-auction mechanisms, performing simulation-based game analysis of heuristic strategies applied to each. In follow-on work, this team showed how to derive new strategies through genetic search over a parametric strategy space, optimizing performance against the equilibrium derived from an empirical game model (Phelps et al., 2006). This was perhaps the first work to automate the outer-loop step of empirical game generation (Figure 1). They further built on these ideas to consider basin size under an evolutionary dynamic as a fitness measure, and proposed an algorithm for extending the strategy space by repeated iteration (Phelps et al., 2010).

The methodology was given the name "EGTA" and systematically developed in a program of sustained research at the University of Michigan, shortly following the Walsh et al. (2002) paper. This began with a study of heuristic strategies for simultaneous ascending auctions (Wellman et al., 2003), which derived constrained equilibria for empirical games over selected parametric instances. A series of PhD dissertations over the next 15 years advanced the methodology in a variety of directions (Cassell, 2014; Jordan, 2010; Reeves, 2005; Schvartzman, 2009; Vorobeychik, 2008; Wiedenbeck, 2015; Wright, 2018).

A significant thread of EGTA work from this group was driven by the Trading Agent Competition (TAC) series of market games, posed as challenges to the research community (Jordan et al., 2010; Niu et al., 2008; Sadeh et al., 2003; Wellman et al., 2007). In these competitions, AI researchers developed innovative trading strategies for a variety of complex market environments. Since the strategies were developed independently (at least initially) by diverse groups focusing on different approaches, understanding the strategic interactions among them often required careful post-hoc analysis. For example, an empirical game study of strategic procurement in the

4 It was foreshadowed though by the first author's earlier dissertation work (Walsh, 2001), in which he characterized "quasi-equilibria" among a restricted class of heuristic strategies in a supply chain game (Walsh et al., 2000).

TAC supply chain game provided insight about why the 2003 tournament was prone to bouts of ruinous price cutting (Wellman et al., 2005). The same game served for a case study in *empirical mechanism design* (EMD) (Vorobeychik et al., 2006), where an incentive engineer searches over mechanism candidates using EGTA to assess strategic response. A recent extension of EMD methodology by Viqueira et al. (2019) also employed a TAC game to exercise the authors' proposed technique.

A particularly exciting future direction for EGTA is the exploitation of recent AI advances—especially deep reinforcement learning (RL)—for the automated generation of strategy candidates. Schvartzman and Wellman (2009b) demonstrated the use of RL with EGTA in an early study of dynamic trading strategies. This has recently been extended to deep RL methods in work at Google DeepMind (Lanctot et al., 2017). It is quite natural that they might pursue this direction, given that their breakthroughs in computer game-playing demonstrated by AlphaGo Zero (Silver et al., 2017) are essentially products of learning through massive simulation of game play. Indeed, DeepMind's advances in the game of StarCraft II explicitly relied on empirical game analysis over the strategy learning process (Vinyals et al., 2019). The promise of combining of deep RL with EGTA has been further demonstrated in complex security games in very recent work (Wang et al., 2019; Wright et al., 2019).

4 Economic Applications of EGTA

Experience employing EGTA in a variety of economic domains illustrates some of the motivating advantages described above. As described in Section 3, much of the original focus of the methodology was on auctions, motivated by scenarios in electronic commerce. There is an extensive game-theoretic literature on auction theory (Krishna, 2010), but definitive treatments are available only for the most highly stylized of situations. Classic auction theory emphasizes one-shot mechanisms like sealed-bid auctions, perhaps with some generalization to handle multiple units, or two-sided bidding. Dealing with multiple heterogeneous goods, simultaneously or sequentially, typically gets beyond the scope of established theory, as does consideration of highly dynamic mechanisms.

Over the years, EGTA studies succeeded in developing insights and best-known strategies for several auction settings that had proven too complex for analytic auction theory. One example studied in early EGTA work (Reeves et al., 2005) is *simultaneous ascending auctions* (SAAs), in which a number of goods are up for sale at the same time. Bidders may submit offers on any subset, as long as these offers exceed the highest prices bid for the respective goods to that point. As bid withdrawal or cancellation is prohibited, prices only ascend, until no agent wishes to bid further and the last bids on each good prevail. Milgrom (2000) had shown that straightforward (myopic) bidding strategies lead to efficient outcomes in SAAs, as long as valuations exhibit gross substitutes. Such strategies could perform quite poorly, however, in the presence of complements, due to the *exposure problem*. Greenwald and Boyan (2004) explored a range of heuristic bidding strategies to deal with exposure, introducing in particular a variety of ways to employ *price predictions* in SAA bidding. In a comprehensive EGTA study, Wellman et al. (2008) considered strategies from prior literature, as well as new strategies based on price predictions that are *self-confirming*: that is, the predictions are borne out when the agents bid accordingly. This work demonstrated the benefits of self-confirming price prediction in the presence of complements, but also found that strategies focusing on demand reduction were more effective when goods are substitutes.

The idea of self-confirming price prediction was likewise found to be effective in simultaneous one-shot auctions, again through a comprehensive EGTA study (Wellman et al., 2017). Follow-on work extended the price-prediction technique to handle interdependence in good prices (Mayer et al., 2013).

In a *continuous double-auction* (CDA), both buyers and sellers bid, with transactions executed as soon as a new bid matches an existing bid in the *order book* (Friedman, 1993). Static double-auctions have been well characterized (Satterthwaite and Williams, 1993), but dynamic CDAs are notoriously difficult to analyze game-theoretically. As noted in the preceding section, Zhan and Friedman (2007) performed EGTA over heuristic markup strategies in a CDA model. A variety of more sophisticated heuristic CDA strategies have been developed in the agent-based literature, incorporating a range of approaches to adaptation and prediction, for example (Cliff, 2006; Gjerstad and Dickhaut, 1998; Tesauro and Das, 2001; Vytelingum et al., 2008). Schvartzman and Wellman (2009b) conducted an EGTA study over representative strategies from the literature and further employed reinforcement learning to generate new strategies improving on these, at least within the environment simulated in that work.

CDAs are of particular relevance as they represent the generic limit-order mechanism at the core of almost all financial markets. ABM researchers have long employed simulation to deal with the intricacies of realistic financial markets (LeBaron, 2006). EGTA finance studies have addressed a variety of financial market contexts, providing insight on existing theoretical models as well as extending understanding to scenarios beyond those apparently addressable by analytic methods. For example, one recent study employed EGTA to explore market making strategies and identify conditions in which market makers were either beneficial or detrimental to market performance (Wah et al., 2017). Other finance topics addressed by EGTA include asset pricing (Cassell and Wellman, 2012), latency arbitrage (Wah and Wellman, 2016), and market manipulation (Wang and Wellman, 2017).

A recent study of banking regulation provides a particularly direct demonstration of what EGTA brings to simulation-based economic reasoning. Poledna et al. (2014) developed an interesting agent-based model capturing the dynamics of *leverage cycles* as described by Geanakoplos (2010). Stated simply, in this model the relatively aggressive agents gain increasing shares of wealth and influence until their predominance renders the system vulnerable to crisis-like corrections. Poledna et al. used this model to study the effects of Basel-style regulation of bank capitalization, finding somewhat counterintuitively that regulation could actually degrade financial stability, as reflected for example in default rates. Given these intriguing results, Cheng and Wellman (2017) undertook to replicate the model and indeed confirmed the original findings. However, they also observed that the comparison of regulated and non-regulated environments did not take into account any possible strategic response of the agents to the presence or absence of regulation. By making aggressiveness a strategic parameter, the EGTA approach could essentially endogenize that choice. The study generated simulation data for profiles over a range of settings, enabling a more appropriate equilibrium-to-equilibrium comparison of the regulated and non-regulated regimes. The main finding was that regulation reduced aggressiveness, and with this response taken into account did not actually exacerbate instability compared to the unregulated environment.

Discussion

As noted in the introduction, agent-based approaches in economics are often motivated by the goal of avoiding the kind of strict rationality assumptions pervasive in mainstream microeconomic theory. Bringing game-theoretic reasoning to agent-based modeling could be viewed as undercutting this particular motivation. That may be fine for those who embrace rationality-based foundations, but others may wish to retain some of the flexibility afforded by the agent-based approach. The EGTA framework itself can accommodate concerns about excessive rationality assumptions in at least three ways.

1. The construction of input strategies may incorporate behavioral assumptions (e.g., myopic or otherwise bounded behavior) or other limitations on full rationality.

2. Since the game environment is much more complex and may already reflect challenges like agent heterogeneity and dynamic uncertainty, notions of equilibrium are much less stylized than in the limited environment. Although behaving according to exact equilibrium may be even more unrealistic, the approach of striving to identify (approximate) equilibrium may be more defensible.

3. The game model is itself subject to analysis with non-standard solution concepts. For example, accommodating behavioral theories such as quantal response or cognitive hierarchy (Camerer, 2003) in the analysis of induced games is quite straightforward.

High-fidelity simulation offers great promise for extending economic analysis to far more complex and realistic situations than are feasible with strictly analytic methods. Particularly with the increasing availability of large-scale computational resources (e.g., through cloud providers), we might expect that economics would begin to leverage simulation technology in a substantial way, just as such computational methods are routinely employed in the physical sciences and engineering. To date adoption has been slow, likely due to methodological questions and resulting acceptance barriers (Lehtinen and Kuorikoski, 2007). Showing how agent-based simulation is not inherently in conflict with accepted concepts, particularly rationality and game-theoretic reasoning, may go some way to weaken those barriers. There yet remain limitations and challenges for simulation-based game analysis. Ongoing work aims to exercise the methodology and refine its technique, building a new and powerful set of simulation-based tools for economic reasoning.

References

Arthur, W. Brian, John H. Holland, Blake LeBaron, Richard Palmer, and Paul Tayler. 1997. Asset Pricing under Endogenous Expectations in an Artificial Stock Market. In W. Brian Arthur, Steven N. Durlauf, and David A. Lane (eds), *The Economy as an Evolving Complex System II*, Boston: Addison-Wesley. 15-43.

Bloembergen, Daan, Karl Tuyls, Daniel Hennes, and Michael Kaisers. 2015. Evolutionary Dynamics of Multi-Agent Learning: A Survey. *Journal of Artificial Intelligence Research*, *53*: 659-697.

Camerer, Colin F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton: Princeton University Press.

Cassell, Ben-Alexander. 2014. *Scaling Empirical Game-Theoretic Analysis*. Ph.D. thesis, University of Michigan.

Cassell, Ben-Alexander and Michael P. Wellman. 2012. Asset Pricing under Ambiguous Information: An Empirical Game-Theoretic Analysis. *Computational and Mathematical Organization Theory*, *18*: 445-462.

Cervone, Roberto, Stefano Galavotti, and Marco LiCalzi. 2009. Symmetric Equilibria in Double Auctions with Markdown Buyers and Markup Sellers. In Cesáreo Hernández, Marta Posada Calvo, and Adolfo López-Paredes (eds), *Artificial Economics: The Generative Method in Economics*, New York: Springer. 81-92.

Chapman, David A. and Valery Polkovnichenko. 2009. First-Order Risk Aversion, Heterogeneity, and Asset Market Outcomes. *Journal of Finance*, *64*(4): 1863-1887.

Cheng, Frank and Michael P. Wellman. 2017. Accounting for Strategic Response in an Agent-Based Model of Financial Regulation. In *Eighteenth ACM Conference on Economics and Computation*. Cambridge, 187-203.

Cliff, Dave. 2006. Evolutionary Optimization of ZIP60: A Controlled Explosion in Hyperspace. In *Eighth Conference on Genetic and Evolutionary Computation*. Seattle, 1621-1628.

Epstein, Joshua M. and Robert Axtell. 1996. *Growing Artificial Societies: Social Science from the Bottom Up*. Cambridge: MIT Press.

Fagiolo, Giorgio and Andrea Roventini. 2012. Macroeconomic Policy in DSGE and Agent-Based Models. *Revue de l'OFCE*, *124*: 67-116.

Fearnley, John, Martin Gairing, Paul Goldberg, and Rahul Savani. 2015. Learning Equilibria of Games via Payoff Queries. *Journal of Machine Learning Research*, *16*: 1305-1344.

Friedman, Daniel. 1993. The Double Auction Market Institution: A Survey. In Daniel Friedman and John Rust (eds), *The Double Auction Market: Institutions, Theories, and Evidence*, Boston: Addison-Wesley. 3-25.

Geanakoplos, John. 2010. The Leverage Cycle. In Daron Acemoglu, Kenneth Rogoff, and Michael Woodford (eds), *NBER Macroeconomics Annual*, Chicago: University of Chicago Press, volume 24. 1-65.

Gilbert, Nigel. 2007. *Agent-Based Models*. New York: Sage.

Gjerstad, Steven and John Dickhaut. 1998. Price Formation in Double Auctions. *Games and Economic Behavior*, *22*(1): 1-29.

Greenwald, Amy and Justin Boyan. 2004. Bidding under Uncertainty: Theory and Experiments. In *Twentieth Conference on Uncertainty in Artificial Intelligence*. Banff, 209-216.

Greenwald, Amy R. and Jeffrey O. Kephart. 1999. Shopbots and Pricebots. In *Sixteenth International Joint Conference on Artificial Intelligence*. Stockholm, 506-511.

Greenwald, Amy R., Jeffrey O. Kephart, and Gerald J. Tesauro. 1999. Strategic Pricebot Dynamics. In *ACM Conference on Electronic Commerce*. Denver, 58-67.

Honorio, Jean and Luis Ortiz. 2015. Learning the Structure and Parameters of Large-Population Graphical Games from Behavioral Data. *Journal of Machine Learning Research*, *16*: 1157-1210.

Jecmen, Steven, Arunesh Sinha, Zun Li, and Long Tran-Thanh. 2020. Bounding Regret in Empirical Games. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*. New York.

Jordan, Patrick R. 2010. *Practical Strategic Reasoning with Applications in Market Games*. Ph.D. thesis, University of Michigan.

Jordan, Patrick R., Yevgeniy Vorobeychik, and Michael P. Wellman. 2008. Searching for Approximate Equilibria in Empirical Games. In *Seventh International Conference on Autonomous Agents and Multi-Agent Systems*. Estoril, 1063-1070.

Jordan, Patrick R., Michael P. Wellman, and Guha Balakrishnan. 2010. Strategy and Mechanism Lessons from the First Ad Auctions Trading Agent Competition. In *Eleventh ACM Conference on Electronic Commerce*. Cambridge, 287-296.

Krishna, Vijay. 2010. *Auction Theory, Second Edition*. Cambridge: Academic Press.

Lanctot, Marc, Vinicius Zambaldi, Audrūnas Gruslys, Angeliki Lazaridou, Karl Tuyls, Julien Pérolat, David Silver, and Thore Graepel. 2017. A Unified Game-Theoretic Approach to Multiagent Reinforcement Learning. In *Thirty-First Annual Conference on Neural Information Processing Systems*. Long Beach, 4193–4206.

LeBaron, Blake. 2006. Agent-Based Computational Finance. In Leigh Tesfatsion and Kenneth L. Judd (eds), *Handbook of Computational Economics*, Amsterdam: Elsevier. 1187-1233.

Lehtinen, Aki and Jaakko Kuorikoski. 2007. Computing the Perfect Model: Why do Economists Shun Simulation? *Philosophy of Science*, *74*(3): 304-329.

Li, Zun and Michael P. Wellman. 2020. Structure Learning for Approximate Solution of Many-Player Games. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*. New York, 2119-2127.

Mayer, Brandon A., Eric Sodomka, Amy Greenwald, and Michael P. Wellman. 2013. Accounting for Price Dependencies in Simultaneous Sealed-Bid Auctions. In *Fourteenth ACM Conference on Electronic Commerce*. Philadelphia, 679-696.

Milgrom, Paul. 2000. Putting Auction Theory to Work: The Simultaneous Ascending Auction. *Journal of Political Economy*, *108*(2): 245-272.

Miller, John H. and Scott E. Page. 2007. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton: Princeton University Press.

Nelson, Richard R. and Sidney G. Winter. 1982. *An Evolutionary Theory of Economic Change*. Cambridge: Harvard University Press.

Niu, Jinzhong, Kai Cai, Simon Parsons, Enrico Gerding, and Peter McBurney. 2008. Characterizing Effective Auction Mechanisms: Insights from the 2007 TAC Market Design Competition. In *Seventh International Conference on Autonomous Agents and Multi-Agent Systems*. Estoril, 1079-1086.

Phelps, S., M. Marcinkiewicz, S. Parsons, and P. McBurney. 2006. A Novel Method for Automatic Strategy Acquisition in N-Player Non-Zero-Sum Games. In *Fifth International Joint Conference on Autonomous Agents and Multi-Agent Systems*. Hakodate, 705-712.

Phelps, Steve, Peter McBurney, and Simon Parsons. 2010. A Novel Method for Automatic Strategy Acquisition and its Application to a Double-Auction Market Game. *IEEE Transactions on Systems, Man, and Cybernetics: Part B*, *40*: 668-674.

Phelps, Steve, Simon Parsons, and Peter McBurney. 2005. An Evolutionary Game-Theoretic Comparison of Two Double-Auction Market Designs. In Peyman Faratin and Juan A. Rodríguez-Aguilar (eds), *Agent-Mediated Electronic Commerce VI: Theories for and Engineering of Distributed Mechanisms and Systems*, New York: Springer. 101-114.

Poledna, Sebastian, Stefan Thurner, J. Doyne Farmer, and John Geanakoplos. 2014. Leverage-Induced Systemic Risk Under Basle II and other Credit Risk Policies. *Journal of Banking and Finance*, *42*: 199-212.

Reeves, Daniel M. 2005. *Generating Trading Agent Strategies: Analytic and Empirical Methods for Infinite and Large Games*. Ph.D. thesis, University of Michigan.

Reeves, Daniel M., Michael P. Wellman, Jeffrey K. MacKie-Mason, and Anna Osepayshvili. 2005. Exploring Bidding Strategies for Market-Based Scheduling. *Decision Support Systems*, *39*(1): 67-85.

Richiardi, Matteo. 2018. Agent-Based Computational Economics: What, Why, When. In Domenico Delli Gatti, Giorgio Fagiolo, Mauro Gallegati, Matteo Richiardi, and Alberto Russo (eds), *Agent-Based Models in Economics: A Toolkit*, Cambridge: Cambridge University Press. 10-32.

Richiardi, Matteo G. 2012. Agent-Based Computational Economics: A Short Introduction. *Knowledge Engineering Review*, *27*(2): 137-149.

Sadeh, Norman, Raghu Arunachalam, Joakim Eriksson, Niclas Finne, and Sverker Janson. 2003. TAC-03: A Supply-Chain Trading Competition. *AI Magazine*, *24*(1): 92-94.

Satterthwaite, Mark A. and Steven R. Williams. 1993. The Bayesian Theory of the k-Double Auction. In Daniel Friedman and John Rust (eds), *The Double Auction Market: Institutions, Theories, and Evidence*, Boston: Addison-Wesley. 99-123.

Schinckus, Christophe. 2019. Agent-Based Modelling and Economic Complexity: A Diversified Perspective. *Journal of Asian Business and Economic Studies*, *26*(2): 170-188.

Schuster, P. and K. Sigmund. 1983. Replicator Dynamics. *Journal of Theoretical Biology*, *100*(3): 533-538.

Schvartzman, L. Julian. 2009. *Stronger Bidding Strategies through Empirical Game-Theoretic Analysis and Reinforcement Learning*. Ph.D. thesis, University of Michigan.

Schvartzman, L. Julian and Michael P. Wellman. 2009a. Exploring Large Strategy Spaces in Empirical Game Modeling. In *AAMAS-09 Workshop on Agent-Mediated Electronic Commerce*. Budapest.

Schvartzman, L. Julian and Michael P. Wellman. 2009b. Stronger CDA Strategies through Empirical Game-Theoretic Analysis and Reinforcement Learning. In *Eighth International Conference on Autonomous Agents and Multi-Agent Systems*. Budapest, 249-256.

Silver, David, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. 2017. Mastering the Game of Go without Human Knowledge. *Nature*, *550*: 354-359.

Sokota, Samuel, Caleb Ho, and Bryce Wiedenbeck. 2019. Learning Deviation Payoffs in Simulation-Based Games. In *Thirty-Third AAAI Conference on Artificial Intelligence*. Honolulu, 1266-1273.

Sureka, Ashish and Peter R. Wurman. 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 Multi-Agent Systems*. Utrecht, 1023-1029.

Taylor, P. and L. Jonker. 1978. Evolutionary Stable Strategies and Game Dynamics. *Mathematical Biosciences*, *40*(1): 145-156.

Tesauro, Gerald and Rajarshi Das. 2001. High-Performance Bidding Agents for the Continuous Double Auction. In *Third ACM Conference on Electronic Commerce*. Tampa, 206-209.

Tesfatsion, Leigh. 2006. Agent-Based Computational Economics: A Constructive Approach to Economic Theory. In Leigh Tesfatsion and Kenneth L. Judd (eds), *Handbook of Computational Economics*, Amsterdam: Elsevier. 831-877.

Tuyls, Karl and Simon Parsons. 2007. What Evolutionary Game Theory Tells Us About Multiagent Learning. *Artificial Intelligence*, *171*(7): 406-416.

Tuyls, Karl, Julien Perolat, Marc Lanctot, Edward Hughes, Richard Everett, Joel Z. Leibo, Csaba Szepesvári, and Thore Graepel. 2020. Bounds and Dynamics for Empirical Game-Theoretic Analysis. *Autonomous Agents and Multi-Agent Systems*, *34*(7).

Vinyals, Oriol, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh, Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, Laurent Sifre, Trevor Cai, John P. Agapiou, Max Jaderberg, Alexander S. Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dalibard, David Budden, Yury Sulsky, James Molloy, Tom L. Paine, Caglar Gulcehre, Ziyu Wang, Tobias Pfaff, Yuhuai Wu, Roman Ring, Dani Yogatama, Dario Wünsch, Katrina McKinney, Oliver Smith, Tom Schaul, Timothy Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps, and David Silver. 2019. Grandmaster Level in StarCraft II using Multi-Agent Reinforcement Learning. *Nature*, *575*: 350-354.

Viqueira, Enrique Areyan, Yasser Mohammed, Cyrus Cousins, and Amy Greenwald. 2019. Empirical Mechanism Design: Designing Mechanisms from Data. In *Thirty-Fifth Conference on Uncertainty in Artificial Intelligence*. Tel Aviv.

Vorobeychik, Yevgeniy. 2008. *Mechanism Design and Analysis Using Simulation-Based Game Models*. Ph.D. thesis, University of Michigan.

Vorobeychik, Yevgeniy. 2010. Probabilistic Analysis of Simulation-Based Games. *ACM Transactions on Modeling and Computer Simulation*, *20*(3): 1-25.

Vorobeychik, Yevgeniy, Christopher Kiekintveld, and Michael P. Wellman. 2006. Empirical Mechanism Design: Methods, with Application to a Supply Chain Scenario. In *Seventh ACM Conference on Electronic Commerce*. Ann Arbor, 306-315.

Vorobeychik, Yevgeniy, Michael P. Wellman, and Satinder Singh. 2007. Learning Payoff Functions in Infinite Games. *Machine Learning*, *67*: 145-168.

Vytelingum, Perukrishnen, Dave Cliff, and Nicholas R. Jennings. 2008. Strategic Bidding in Continuous Double Auctions. *Artificial Intelligence*, *172*(14): 1700-1729.

Wah, Elaine and Michael P. Wellman. 2016. Latency Arbitrage in Fragmented Markets: A Strategic Agent-Based Analysis. *Algorithmic Finance*, *5*(3): 69-93.

Wah, Elaine, Mason Wright, and Michael P. Wellman. 2017. Welfare Effects of Market Making in Continuous Double Auctions. *Journal of Artificial Intelligence Research*, *59*: 613-650.

Walsh, William E. 2001. *Market Protocols for Decentralized Supply Chain Formation*. Ph.D. thesis, University of Michigan.

Walsh, William E., Rajarshi Das, Gerald Tesauro, and Jeffrey O. Kephart. 2002. Analyzing Complex Strategic Interactions in Multi-Agent Systems.

In *AAAI-02 Workshop on Game-Theoretic and Decision-Theoretic Agents*. Edmonton.

Walsh, William E., David Parkes, and Rajarshi Das. 2003. Choosing Samples to Compute Heuristic-Strategy Nash Equilibrium. In *AAMAS-03 Workshop on Agent-Mediated Electronic Commerce*. Melbourne, 109-123.

Walsh, William E., Michael P. Wellman, and Fredrik Ygge. 2000. Combinatorial Auctions for Supply Chain Formation. In *Second ACM Conference on Electronic Commerce*. Minneapolis, 260-269.

Wang, Xintong and Michael P. Wellman. 2017. Spoofing the Limit Order Book: An Agent-Based Model. In *Sixteenth International Conference on Autonomous Agents and Multi-Agent Systems*. São Paulo, 651-659.

Wang, Yufei, Zheyuan Ryan Shi, Lantao Yu, Yi Wu, Rohit Singh, Lucas Joppa, and Fei Fang. 2019. Deep Reinforcement Learning for Green Security Games with Real-Time Information. In *Thirty-Third AAAI Conference on Artificial Intelligence*.

Weiss, Gerhard, Lars Braubach, and Paolo Giorgini. 2010. Intelligent Agents. In Hossein Bidgoli (ed.), *Handbook of Technology Management*, Hoboken: Wiley. 360-372.

Wellman, Michael P. 2006. Methods for Empirical Game-Theoretic Analysis (Extended Abstract). In *Twenty-First AAAI Conference on Artificial Intelligence*. Boston, 1552-1555.

Wellman, Michael P. 2016. Putting the Agent in Agent-Based Modeling. *Autonomous Agents and Multi-Agent Systems*, *30*: 1175-1189.

Wellman, Michael P., Joshua Estelle, Satinder Singh, Yevgeniy Vorobeychik, Christopher Kiekintveld, and Vishal Soni. 2005. Strategic Interactions in a Supply Chain Game. *Computational Intelligence*, *21*: 1-26.

Wellman, Michael P., Amy Greenwald, and Peter Stone. 2007. *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*. Cambridge: MIT Press.

Wellman, Michael P., Tae Hyung Kim, and Quang Duong. 2013. Analyzing Incentives for Protocol Compliance in Complex Domains: A Case Study of Introduction-Based Routing. In *Twelfth Workshop on the Economics of Information Security*. Washington.

Wellman, Michael P., Jeffrey K. MacKie-Mason, Daniel M. Reeves, and Sowmya Swaminathan. 2003. Exploring Bidding Strategies for Market-Based Scheduling. In *Fourth ACM Conference on Electronic Commerce*. San Diego, 115-124.

Wellman, Michael P., Anna Osepayshvili, Jeffrey K. MacKie-Mason, and Daniel M. Reeves. 2008. Bidding Strategies for Simultaneous Ascending Auctions. *B. E. Journal of Theoretical Economics (Topics)*, *8*(1).

Wellman, Michael P., Daniel M. Reeves, Kevin M. Lochner, and Rahul Suri. 2006. Searching for Walverine 2005. In *Agent-Mediated Electronic Commerce: Designing Trading Agents and Mechanisms*, number 3937 in Lecture Notes in Artificial Intelligence, New York: Springer. 157-170.

Wellman, Michael P., Eric Sodomka, and Amy Greenwald. 2017. Self-Confirming Price-Prediction Strategies for Simultaneous One-Shot Auctions. *Games and Economic Behavior*, *201*: 339-372.

Wiedenbeck, Bryce. 2015. *Approximate Analysis of Large Simulation-Based Games*. Ph.D. thesis, University of Michigan.

Wiedenbeck, Bryce, Ben-Alexander Cassell, and Michael P. Wellman. 2014. Bootstrap Techniques for Empirical Games. In *Thirteenth International Conference on Autonomous Agents and Multi-Agent Systems*. Paris, 597-604.

Wiedenbeck, Bryce, Fengjun Yang, and Michael P. Wellman. 2018. A regression Approach for Modeling Games with Many Symmetric Players. In *Thirty-Second AAAI Conference on Artificial Intelligence*. New Orleans, 1266-1273.

Wright, Mason. 2018. *Stable Profiles in Simulation-Based Games via Reinforcement Learning and Statistics*. Ph.D. thesis, University of Michigan.

Wright, Mason, Yongzhao Wang, and Michael P. Wellman. 2019. Iterated Deep Reinforcement Learning in Games: History-Aware Training for Improved Stability. In *Twentieth ACM Conference on Economics and Computation*. Phoenix, 617-636.

Zhan, Wenjie and Daniel Friedman. 2007. Markups in Double Auction Markets. *Journal of Economic Dynamics and Control*, *31*(9): 2984-3005.