

Evolutionary Mechanism Design: A Review

Steve Phelps · Peter McBurney · Simon Parsons

Received: date / Accepted: date

Abstract The advent of large-scale distributed systems poses unique engineering challenges. In open systems such as the internet it is not possible to prescribe the behaviour of all of the components of the system in advance. Rather, we attempt to design infrastructure, such as network protocols, in such a way that the overall system is robust despite the fact that numerous arbitrary, non-certified, third-party components can connect to our system. Economists have long understood this issue, since it is analogous to the design of the rules governing auctions and other marketplaces, in which we attempt to achieve socially-desirable outcomes despite the impossibility of prescribing the exact behaviour of the market participants, who may attempt to subvert the market for their own personal gain. This field is known as "mechanism design": the science of designing rules of a game to achieve a specific outcome, even though each participant may be self-interested. Although it originated in economics, mechanism design has become an important foundation of multi-agent systems (MAS) research. In a traditional mechanism design problem, analytical methods are used to prove that agents' game-theoretically optimal strategies lead to socially desirable outcomes. In many scenarios, traditional mechanism design and auction theory yield clear-cut results; however, there are many situations in which the underlying assumptions of the theory are violated due to the messiness of the real-world. In this paper we review alternative approaches

We are grateful for financial support received from the UK EPSRC through the *Market-Based Control of Complex Computational Systems* Project (GR/T10657/01), and from the US National Science Foundation under grant NSF IIS-0329037 *Tools and Techniques for Automated Mechanism Design*.

S. Phelps

Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, Wivenhoe Park, Colchester CO4 3SQ, United Kingdom

Tel.: +44-(0)1206-874414

Fax: +44 (0) 1206 874122

E-mail: sphelps@essex.ac.uk

P. McBurney

Department of Computer Science, University of Liverpool, Liverpool L69 3BX UK

E-mail: mcburney@liverpool.ac.uk

S. Parsons

Department of Computer and Information Science, Brooklyn College, City University of New York, 2900 Bedford Avenue, Brooklyn, 11210 NY USA

E-mail: parsons@sci.brooklyn.cuny.edu

to mechanism design which treat it as an *engineering* problem and bring to bear engineering design principles, viz.: iterative step-wise refinement of solutions, and satisficing instead of optimization in the face of intractable complexity. We categorize these approaches under the banner of *evolutionary* mechanism design.

1 Introduction

Much work in the design of multi-agent systems (MAS) has focused on the design and engineering of individual agents; for example, the problems of designing and implementing effective trading strategies for agents participating in e-commerce market places, or the design of effective learning algorithms for adaptive agents. However, increasingly attention is being turned to the design of the infrastructure, or the environment, underlying the interactions between individual agents in a MAS; for example, the problem of designing rules governing the operation of an e-commerce market institution, or the design of interaction protocols governing agent argumentation. The justification for the latter approach is that often as MAS designers we are responsible for engineering *open* systems, in which we do not have control over the exact behavior of the agents connecting to our system; these agents are, after all, autonomous. Rather, we build a set of standards and protocols that define a framework within which our agents are free to interact, and if we have designed this framework to be robust, the system as a whole will exhibit our desired design properties despite the fact that it consists of possibly millions of autonomous agents interacting with each other in ways we have not prescribed in advance.

Economists have studied similar design problems in the context of *auction theory* (Klemperer, 2004a; Krishna, 2002) and *mechanism design* (Jackson, 2003; Varian, 1995). In a mechanism design problem, the task of the designer is to choose the rules of the auction in such a way that the designer's objectives are met when agents play their optimal strategies. One of the main difficulties in solving this problem is computing the optimal strategies, as the best strategy to play depends on what strategies are being played by other agents; the number of agents can vary significantly, and the strategy space can be very large. The standard technique is to view each possible set of auction rules as defining a particular game, and then to use game theory to "solve" this game by finding the set of strategies comprising a *Nash equilibrium* of the game — the set of strategies that are best responses to each other. For many scenarios, especially for single-sided auctions comprising a single seller and multiple buyers, auction theory and mechanism design yield clear-cut results. However, in the general case the problem is analytically intractable, especially when it comes to analysing *double-sided* auctions, also known as exchanges, in which we allow multiple sellers as well as multiple-buyers. In the next section we shall describe our motivation for studying double-sided auctions.

A double-auction mechanism is a generalization of an auction in which there are multiple sellers as well as multiple buyers, and both buyers and sellers are allowed to exchange offers simultaneously. Since double-auctions allow dynamic pricing on both the supply side and the demand side of the marketplace, their study is of great importance, both to theoretical economists (Klemperer, 2004b), and those seeking to implement real-world market places. On the one hand, economists who are interested in theories of price formation in idealized models of general markets have often turned to exchange-like models such as Walrasian tâtonnement, to describe and understand the price-formation process (Bryant, 2000), and on the other hand, variants of the double-auction are used in large real-world exchanges to trade

commodities in marketplaces where supply and demand fluctuate rapidly, such as markets for stocks, futures, and their derivatives.

However, the models of exchanges traditionally used by economists in general equilibrium theory are often simplified for the purposes of analytical tractability to such an extent that they are of scant relevance to the designers of real-world exchanges, and even, it is sometimes argued, of scant relevance to the theoretical modelling of markets (Fisher, 1983). For example, one important simplification often made is that the number of agents participating in a market is very large; this simplification allows relative market power and consequent *strategic effects* to be ignored. However, in many real-world marketplaces, such as deregulated wholesale electricity markets, there may be relatively few competitors on one or both sides of the market. With small numbers of participants, general equilibrium models break down (Medio and Gallo, 1992, p. 10) because they fail to allow for market power, and the potential gains of strategic behavior, of participants.

1.1 Auction Theory & Mechanism Design

Auction theory can be thought of as an alternative approach to general equilibrium theory, in which we build a more sophisticated micro-model of the marketplace, and we use game-theoretic techniques to analyse the rational behavior of *individual agents* faced with different pricing institutions. Whereas neoclassical equilibrium theory often abstracts away from the details of individual agents, game-theoretic models allow economists to build sophisticated micro-models of individual agents' reasoning and preferences. In many scenarios, especially in analyzing single-sided monopoly markets, these models have been spectacularly successful to the extent where they have been directly applied to the design of real-world auctions for high-value government and corporate assets (Klemperer, 2002). However, in other practical scenarios, especially when it comes to analyzing and designing double-sided markets, such as exchanges, there are still a number of problems with the theory, which we shall briefly review.

Auction-theorists typically analyze a proposed market institution by defining a set of design objectives, and then proceed to show that these design objectives are brought about when rational agents follow their best strategies according to a game-theoretic analysis. The task of choosing the rules of the market institution so that these objectives are brought about is called *mechanism design*. The typical design objectives considered by mechanism designers are:

Allocative efficiency: The outcome of using the mechanism should be optimal in some defined sense, for example, the total surplus generated should equal the available surplus in competitive equilibrium.

Budget balance: No outside subsidy inwards or transfers outwards are required for a deal to be reached.

Individual rationality: The expected net benefit to each participant from using the mechanism should be no less than the net benefit of any alternative.

Incentive compatibility: Participants should not be able to gain an advantage from non-truth-telling behavior.

We would like auction mechanisms to satisfy all of these properties. However, this is not possible in many situations. For example, the impossibility result of (Myerson and Satterthwaite, 1983) demonstrates that no *double-sided* auction mechanism can be simultaneously efficient, budget-balanced and individually-rational. Moreover, many of the underpinnings

of the theory assume that designers' interests are restricted to only the aforementioned properties. For example, the revelation principle (Krishna, 2002, p. 62) states that, without loss of generality, we may safely restrict attention to mechanisms in which agents reveal their types truthfully. However, this result does not take into account the potential cost or other practicalities of polling agents for their type information. Once minimizing the cost of revelation is introduced as a design objective, the revelation principle ceases to hold, because there may exist partial-revelation mechanisms with non-truthful equilibria which sacrifice incentive-compatibility for expedience of revelation. This is of more than academic interest, since in real-world electronic exchanges it is rarely possible to poll *all* agents for their valuations before clearing the market; hence the *continuous* double-auction, in which we execute the clearing operation as new offers arrive, thus increasing liquidity (transaction throughput) at the expense of incentive-compatibility.

In designing market places, as with any other engineering problem, we often need to make such trade-offs between different objectives depending on the exact requirements and scenario at hand. We can often satisfactorily solve such multi-objective optimisation problems, provided that we have some kind of quantitative assessment of each objective, yet classical auction-theory provides only a binary yes or no indication of whether each of its limited design objectives is achievable, making it extremely difficult to compare the different trade-offs.

Further complications arise when we attempt to use auction-theory to analyze existing ("legacy") market institutions. Exchanges such as the London Stock Exchange¹ have been in existence far longer than game-theory and auction-theory, thus, unsurprisingly, the original rules of the institution were not necessarily based on sound game-theoretic or auction-theoretic principles. Moreover, it is unrealistic to expect that core financial institutions such as these radically alter their rules overnight in response to the latest fashionable developments in auction-theory or game-theory. Rather, it may be more salient to view financial institutions *evolving* gradually and incrementally in response to a changing environment. Similarly, agents participating in these institutions do not necessarily instantaneously and simultaneously adjust their trading behavior to the theoretical optimum strategy; for example, adoption of a new trading strategy may spread through a population of traders as word of its efficacy diffuses in a manner akin to mimetic evolution.² Thus, we may think of the institutions we see today as the outcome of a *co-evolutionary* adaptation between financial institutions on the one hand, and trading strategies on the other.

The issue of legacy institutions has ramifications for mechanism design; in these contexts the choice of adjustments to the auction rules may be tightly constrained by existing infrastructure, both physical and social; thus it may be necessary to examine the *attainability* of equilibria under the new design given existing strategic behavior in the legacy design. Classical auction theory relies on classical game-theory which in turn says nothing about the *dynamics* of adjustment to equilibrium.

For such applications, we need to turn to models of evolution and learning in strategic environments; models that we collectively categorize under the banner of *evolutionary game theory*. Models of learning and evolution as applied to agents' strategies are not new. Where our approach differs, however, is in the application of models of learning and evolution to the market mechanism itself, a new field we call *evolutionary mechanism design*.

¹ www.londonstockexchange.com

² The adoption by derivatives traders of the Black-Scholes equation for option pricing provides an example (MacKenzie, 2003).

The remainder of this paper is outlined as follows. In Section 2 we review the literature on the application of AI to the modelling of agent interactions in strategic environments such as stock markets and distributed computing systems. In Section 3 we review the emerging literature on the application of AI to the design of the *infrastructure* underlying these interactions, i.e. the mechanism design problem, with an emphasis on own work in this area. Finally we conclude in Section 4 and discuss future research directions.

2 Agent-based Economics

It has long been understood that Artificial Intelligence (AI) has strong roots in economics (Russell and Norvig, 2003, p. 9); whilst the latter is traditionally concerned with idealized models of agents interacting in realistically complex environments, the former has placed more emphasis on realistically complex agents interacting in idealized environments. Indeed, one of the pioneers of AI, Herbert Simon was originally motivated in much of his AI research by attempts to build more complex models of agents' behaviour in economic environments (see, for example Clarkson and Simon (1960)).

Whilst the broad relationships between the two disciplines were generally understood from the inception of AI, it was not until the late twentieth century and the birth of the Multi-Agent Systems (MAS) discipline that highly specialised theories and concepts were imported from economics into AI. Boutilier et al (1997) were amongst the first to clearly articulate the specific relationships between economics and AI. The particular significance of mechanism design in the context of multi-agent systems was first discussed by Rosenschein and Zlotkin (1994) and Varian (1995), as summarised by Wellman:

“Within economics, the problem of synthesizing an interaction protocol via which rational agents achieve a socially desirable end is called mechanism design. This is exactly the problem we face in designing distributed software systems, which suggests an opportunity to exploit existing economic ideas.” (Wellman, 1995)

More recently the theme of incentive engineering has been taken up in the wider computer science community in contexts as diverse as information security (Anderson, 2001), and computer networking:

“If an artifact (a new congestion control protocol, a new caching scheme, a new routing algorithm, etc.) is demonstrated to have superior performance, this does not necessarily mean that it will be successful. For the artifact to be ‘fit’, there must exist a path leading from the present situation to its prevalence. This path must be paved with incentives that will motivate all kinds of diverse agents to adopt it, implement it, use it, interface with it or just tolerate it. In the absence of such a path, the most clever, fast and reliable piece of software may stay just that. All design problems are now mechanism design problems.” (Papadimitriou, 2001)

2.1 The Double Auction

Our work focuses specifically on a particular class of economic mechanism — the double auction. As discussed in the previous section, the double auction has come to be recognized as an important *benchmark problem*, in both economics and multi-agent systems. In particular, a landmark workshop held in Santa Fe motivated much contemporary research in this

area by highlighting the difficulty of agents' decision problems in non-idealized variants of this type of marketplace, and the Santa Fe double-auction tournament was one of the first studies which used advanced agent-based simulation in order to explore the properties of this mechanism (Friedman and Rust, 1991). To this day the double-auction still represents an important benchmark problem by simultaneously admitting of precise representations whilst stretching the bounds of both analytical tractability and computational brute-force. In the following we will review analytical and computational approaches to the agents' decision problem (traditionally the focus of AI), and the mechanism-design problem (traditionally the focus of economics) in turn.

2.1.1 Analytical approaches

The core of the analytic approach to agents' decision problems is based around the theory of n-player non-zero-sum games as formulated by Nash (1950). Nash's insight was that in any interaction of preference-maximising agents whose outcome depends on the joint set of actions — that is, a game — any given agent has a theoretical *best response* to the actions chosen by the other agents. By applying this reasoning recursively we arrive at the concept of a Nash equilibrium; a situation in which every agent chooses actions that are best-responses to the best-responses of other agents. Nash proved that provided players have the option of randomising over their choice of strategies, then *every* n-player game possesses at least one equilibrium solution, thus providing a powerful theoretical framework not only for optimizing one's strategy in such an interaction (choosing a best-response), but also in predicting a likely combination of joint actions (Nash equilibrium). Many refinements have since been made to Nash's theory, some of the most important being the concept of a Bayesian-Nash equilibrium (BNE) (Harsanyi, 1967), which deals with situations where payoffs are dependent on some private unobservable properties of an agent — the agent's *type* (for example, the particular cards that an agent holds in a game of poker), and the theory of evolutionary games (Maynard-Smith, 1982; Gintis, 2000) which overlays a dynamic model of gradual strategy-adjustment on top of the static equilibria of Nash's original formulation.

Game-theory provides a very powerful *general* framework for solving agent interactions in theory, but it was William Vickrey (Vickrey, 1961, 1962) who first saw the fundamental economic significance of auctions and who first applied the theory of games in this area giving birth to modern auction theory, as summarised by Krishna (2002).

Auction theory provides a comprehensive theoretical framework for analysing single sided auctions — that is, auctions with a single seller and multiple buyers. However, double-sided auctions — auctions with multiple sellers as well as multiple buyers — remain something of a theoretical oddity despite their increasing prevalence in economic reality. Vickrey (1961) demonstrated that no double-sided mechanism could simultaneously achieve the incentive compatibility, individual-rationality, budget-balance and efficiency desiderata. Subsequently d'Aspremont and Gérard-Varet (1979) demonstrated the existence of a budget-balanced mechanism that was able to achieve incentive-compatibility in Bayesian-Nash equilibrium³ at the expense of individual rationality. McAfee (1992) provided a formulation of a double-sided single-unit mechanism that admitted of a dominant-strategy game-theoretic solution at the expense of budget-balance, and (Huang et al, 2002) later refined this idea to the multi-unit case. However Myerson and Satterthwaite (1983); Satterthwaite and Williams (1993) were able to extend Vickrey's result and demonstrated that for the case

³ Bayesian-Nash incentive-compatibility merely requires truth-telling as Bayesian-Nash equilibrium of the game, rather than the usual stricter requirement that truth-telling is a dominant-strategy.

of a single buyer and seller there does not exist a mechanism that can simultaneously achieve incentive-compatibility, budget-balance, efficiency and individual-rationality even when the incentive-compatibility criteria is relaxed from dominant-strategy to BNE, and hence there is no double-sided mechanism for achieving all the usual desiderata required by auction theorists in the general case.

Although there is no unequivocal and complete game-theoretic analysis of the double-auction in the general case, that is not to say, however, that double-sided mechanisms do not admit of game-theoretic solutions in specific instances. The first equilibrium analysis for a double auction was that of Chatterjee and Samuelson (1983), in the paper in which they introduced the idea of the k -double auction⁴, which we will discuss in the next chapter, albeit only for the two trader case. In this initial paper, Chatterjee and Samuelson show that there is an equilibrium solution, assuming independent private values.

Considerable work has since been carried out extending this result. First, Williams (1991, 1988) showed the existence of equilibria in the buyer's bid double auction — this is an easier auction to analyse since the dominant strategy for sellers is to bid their true value, thus fixing one side of the auction and, as Rustichini et al (1994) points out, ensuring that the market has a unique equilibrium⁵. The same authors subsequently showed the existence of equilibria in the many-trader version of the k -double auction (Satterthwaite and Williams, 1993), at the same time suggesting that the modified BBDA has no equilibrium. This work was followed by Jackson and Swinkels (2001, 2005), who showed the existence of equilibria, though not monotonic equilibria, under a wide range of conditions. Next, Reny and Perry (2003) showed that monotonic equilibria exist if offers are restricted to discrete values, and Fudenberg et al (2003) showed that this result could be extended to continuous values (which Jackson and Swinkels (2005) argue is “a very useful approximation . . . allowing one . . . to use calculus to characterise equilibria”) provided that the auction was large. Finally, Kadan (2004) showed that an increasing equilibrium exists for just two traders with affiliated values.

2.1.2 Empirical and agent-based approaches

Whilst double-auction mechanisms stretch the bounds of auction-theory by admitting of no unequivocal dominant strategy solution in the general case, the theory of games itself has come under scrutiny as a plausible general-purpose model of the strategic behavior of complex agents (human or otherwise); for example, Goeree and Holt (2001) give an overview of ten simple games where the game-theoretic solution is easily obtainable yet intuitively implausible. This has led to a re-examination of the use of *empirical* methods in economics, whereby experiments are conducted with actual agents trading in a market-institution under laboratory conditions. The agents may be human, in which case the methodology is sometimes called *experimental economics* (see for example Kagel and Roth (1995)), or more generally they may be implemented in the form of a computer-program; Tesfatsion (2002) coined the phrase *agent-based computational economics* (ACE), to describe this approach.

Experimental economics using human agents has the advantage that a large supply of agents are available “off the shelf” so to speak; hence not surprisingly experiments using human agents were among the first ACE investigations of the double-auction market. Smith

⁴ Though not under this name — they refer to the price setting rule as a “bargaining rule”.

⁵ Note that all results for the BBDA, the 1-DA, are symmetric with those for the k -double auction in which the transaction price is determined by the price offered by the highest asking seller that trades, the 0-DA (Rustichini et al, 1994).

(1962) was the first to study the double-auction under laboratory conditions using human-agents, and his results suggested that human subjects were able to extract close to theoretically optimal surplus from the market.

One of the disadvantages of human-based experimental economics compared with agent-based computational economics is that it is not always straightforward to analyze the necessary cognitive mechanisms required to achieve a particular economic outcome. In contrast Gode and Sunder (1993) performed one of the earliest agent-based experiments on the double-auction with the aim of investigating the lower-bounds on the amount of cognitive machinery required to achieve efficient outcomes. They were able to demonstrate that their minimal *zero-intelligence* strategies, implemented in the form of computer programs, were able to achieve highly efficient outcomes, suggesting that the double-auction mechanism was highly robust in the sense that it required minimal rationality on behalf of participants. Their results were not unequivocal, however; Cliff and Bruten (1997) demonstrated that some aspects of Gode and Sunder's results were highly contingent on the particular distribution of agents' valuations that were used in the original experiments, and that a more sophisticated and robust strategy, *zero-intelligence plus* (ZIP) was required in order more accurately fit the behaviour of human subjects under less restrictive assumptions.

This was not the end of the story, though, since when analysing a market mechanism ideally we want to demonstrate the existence of a *dominant* strategy, and that design objectives such as high-efficiency outcomes are the result of agents adopting this particular strategy. For example, in many single-sided auctions one of the desiderata usually considered is *incentive-compatibility*; the dominant bidding strategy should be to bid truthfully at one's valuation. Unless we can demonstrate that an economic outcome such as high efficiency is the result of agents adopting a dominant strategy, or at the very least an equilibrium strategy profile, we can never be sure that the strategy under which high efficiency is observed will not, at some point, be discarded in favour of an alternative strategy which yields higher payoff for its adopters at the expense of overall social welfare. By analogy, consider the prisoner's dilemma game (Flood, 1952; Bendor and Swistak, 1997; Axelrod, 1997); although the cooperative strategy yields the highest welfare outcome if all agents adopt it, this does not suffice to demonstrate that both agents will adopt the cooperative strategy since there will always be a temptation to choose the defection strategy.

Thus there have been numerous attempts to craft agent-based trading strategies for double-auctions that are able to out-compete other strategies: Preist and van Tol (2003) devised a variant of Cliff's ZIP strategy that was able to trade in persistent-shout auctions; Gjerstad and Dickhaut (2001) (GD) introduced a trading strategy that estimates the probability of a bid being accepted as a function of bid price based on an analysis of historical market data, and then bids to maximise expected profit; this approach was refined in later work by Tesauro and Bredin (2002) to incorporate a forecasting element; Todd Kaplan's (Friedman and Rust, 1991) entry into the Santa Fe tournament was one of the first documented double-auction *sniping* strategies, which waits either until the market is in its favour, or the auction is about to close before submitting a bid; Tesauro and Das (2001) introduced variants of the GD and ZIP strategies that were able to trade in continuous-time environments; Nicolaisen et al (2001) used a trading strategy based on a reinforcement-learning model of human game playing (Erev and Roth, 1998) to analyse a simulated electricity market; and Hsu and Soo (2001) analysed the performance of a strategy based on the q-learning algorithm (Watkins and Dayan, 1992). Variations on these and other strategies have been pitted against each other in several public tournaments designed to elicit new strategy designs from the ACE community (Greenwald and Stone, 2001; Wellman et al, 2001). Although some of these strategies have advantages over others in certain situations, and there are pros and cons to

each, there is evidence to suggest that none of them are *dominant* over the others (Walsh et al, 2002), even putting aside the problem of demonstrating that any are dominant over the entire space of possible strategies.

2.1.3 Evolutionary search

Much of the work cited in the previous section focussed on showing that particular strategies yield high payoff if deployed in a market in which all agents adopt the same strategy homogeneously. However, if we have reason to believe that none of the strategies from the previous section are dominant over the others when they interact with each other in the same marketplace, we have no reason to believe that any *single* one of them will come to be used in a real market. Hence if we simply compute market outcomes by running experiments in which we equip agents homogeneously with the same non-dominant strategy, we are not necessarily nearer to understanding the economic properties of the double-auction.

Of course, it may be the case that a single dominant strategy simply does not exist for the double-auction game; instead, some *mixture* of these, or yet to be discovered strategies, might constitute a Nash *equilibrium*. That is, even though no single strategy is “optimal” in the sense that it is dominant over the others, some mix of these or other strategies might constitute best-responses to each other. If this were the case and our market were populated by such a mix of strategies, we might expect that such a state of affairs would persist in reality, since by definition if the agents were to change their strategy they would be worse off. Therefore the agents themselves would have an incentive to maintain the status quo; and thus the components of the system would tend to naturally drive the system back towards such an equilibrium. Thus if we evaluate the properties of the mechanism when it is in these equilibrium states, we might expect that our predictions for variables such as market-efficiency will be accurate for some reasonable duration, and if our design objectives are maximised in these equilibria we will have shown that our mechanism is homeostatic.

In order to assess whether or not there are mixtures of strategies constituting equilibria, it is necessary to systematically evaluate the strategic *interaction* between the known strategies, as well as the space of yet to be considered strategies. Since this search-space is very large, exhaustive search is unfeasible. This has led researchers to turn to heuristic methods such as evolutionary search as a possible methods for studying the interaction between different double-auction strategies by systematically sampling the search space, e.g.: Cliff (1998) used evolutionary search to explore the parameter space of the ZIP strategy, and Andrews and Prager (1994) used Koza’s genetic programming technique (Koza, 1993) to search for a best-response to a uniform mixed-strategy of the Santa Fe tournament entries. *Co*-evolutionary algorithms (Hillis, 1992; Angeline and Pollack, 1993; Pollack and Blair, 1998) are highly promising in this respect. In a *co*-evolutionary search the fitness of individuals in the population is evaluated relative to one another in joint interactions (similarly to payoffs in a strategic game), and it is suggested that in certain circumstances the converged population is an approximate Nash solution to the underlying game; that is, the stable states, or equilibria, of the *co*-evolutionary process are related to the game-theoretic equilibria. Price (1997) and Dawid (1999) used *co*-evolutionary search to explore convergence to equilibrium states in the double-auction.

However, there are many caveats to interpreting the equilibrium states of standard *co*-evolutionary algorithms as approximations of game-theoretic equilibria, as discussed in detail by Sevan Ficici (Ficici and Pollack, 2000, 1998). This has led to a number of refinements to standard *co*-evolutionary algorithms by incorporating game-theoretic concepts directly into the *co*-evolutionary algorithm itself (Noble and Watson, 2001; Ficici and Pollack,

2003; Ficici, 2004); the use of heuristic search (evolutionary or otherwise) to find approximate best-response or equilibrium strategies is a topic that we shall return to in Section 3.3.

2.2 A hybrid approach: empirical game-theory

The various caveats discussed above with the game-theoretic, agent-based and evolutionary approaches, as used in isolation, have inspired *hybrid* approaches whereby agent-based experimentation is used to build an approximate game-theoretic representation which is then solved using standard techniques from classical and evolutionary game-theory. This methodology is known as *empirical* game-theory, and it is the principle methodology used in our work (Section 3.2.1). Many studies prior to 2000 had started to take a more principled and systematic approach to studying the interaction between complex strategies in a simulation context. For example Rust, Miller and Palmer systematically studied convergence to equilibrium of the strategies in the original Santa Fe tournament using ideas very similar to evolutionary game-theory (Friedman and Rust, 1991, p. 183–189). These ideas matured within the MAS community, and Wellman and colleagues set this kind of analysis in a rigorous game-theoretic terms: Walsh et al (2002) demonstrated the effectiveness of the technique for several bargaining games, including a double-auction; Walsh et al (2003) introduced a refinement to the technique to concentrate the sampling of simulations on those experiments that were most critical to the equilibrium analysis ; Reeves et al (2005) performed a game-theoretic analysis of strategies in a market-based scheduling scenario and Jordan et al (2007); Wellman et al (2005) used empirical game-theoretic analysis in relation to the trading agent competition and other games (Jordan et al, 2008).

The basic idea of empirical game theory is to obviate many of the tractability problems discussed in previous sections by restricting attention to small representative sample of “heuristic” strategies that are known to be commonly played in a given multi-state game. For many games, unsurprisingly none of the strategies commonly in use is dominant over the others. Given the lack of a dominant strategy, it is then natural to ask if there are mixtures of these “pure” strategies that constitute game-theoretic equilibria.

For small numbers of players and heuristic strategies, we can construct a relatively small normal-form payoff matrix which is amenable to game-theoretic analysis. This *heuristic* payoff matrix is calibrated by running many iterations of the game; variations in payoffs due to different player-types (eg private valuations) or stochastic environmental factors (e.g. PRNG seed) are averaged over many samples of type information resulting in a single mean payoff to each player for each entry in the payoff matrix. Players’ types are assumed to be drawn independently from the same distribution, and an agent’s choice of strategy is assumed to be independent of its type, which allows the payoff matrix to be further compressed, since we simply need to specify the number of agents playing each strategy to determine the expected payoff to each agent. Thus for a game with j strategies, we represent entries in the heuristic payoff matrix as vectors of the form

$$\mathbf{p} = (p_1, \dots, p_j)$$

where p_i specifies the number of agents who are playing the i^{th} strategy. Each entry $p \in P$ is mapped onto an outcome vector $q \in Q$ of the form

$$\mathbf{q} = (q_1, \dots, q_j)$$

where q_i specifies the expected payoff to the i^{th} strategy. For a game with n agents, the number of entries in the payoff matrix is given by

$$s = \frac{(n + j - 1)!}{n!(j - 1)!} \quad (1)$$

For small n and small j this results in payoff matrices of manageable size; for $j = 3$ and $n = 6, 8,$ and 10 we have $s = 28, 45,$ and 66 respectively.

Once the payoff matrix has been computed we can subject it to a rigorous game-theoretic analysis, search for Nash equilibria solutions, and apply different models of learning and evolution, such as the replicator dynamics model (Maynard-Smith, 1982), in order to analyse the dynamics of adjustment to equilibrium.

The equilibria solutions that are thus obtained are not rigorous Nash equilibria for the full multi-state game; there is always the possibility that an unconsidered strategy could invade the equilibrium. Nevertheless, heuristic-strategy equilibria are sometimes more plausible as models of real-world game playing than those obtained using, for example, a co-evolutionary search, precisely because they *restrict* attention to strategies that are commonly known and are in common use. We can therefore be confident that no commonly known strategy for the game at hand will break our equilibrium, and thus the equilibrium stands at least some chance of persisting in the short term future. In Section 3.2.1 we shall review the application of empirical game theory to the mechanism design problem.

3 Evolutionary mechanism design

Whilst the application of computational techniques to the agent decision problem has a comparatively long tradition, their application to the mechanism-design problem is more recent. The economist Alvin Roth was the first to pose mechanism-design as an *engineering* problem (Roth, 2002), thus paving the way for the application of engineering techniques to mechanism-design. The authors (Phelps et al, 2002b,a) and Cliff (2001) were the first to apply evolutionary search to the double-auction design problem with a view to automating the mechanism-design process. Subsequently Bye (2003) used evolutionary computing to analyze a space of variants to the Vickrey nth-pricing rule in the context of single-sided auctions.

Meanwhile, more generally the research areas of *automated* mechanism design (Conitzer and Sandholm, 2002) and *computational* mechanism design (Dash et al, 2003) were emerging. Conitzer and Sandholm (Conitzer and Sandholm, 2002; Sandholm, 2003) were the first to pose the automated mechanism-design problem in rigorous theoretical terms and to analyze the algorithmic complexity of the problem. This has also led to ideas in adaptive mechanism design in which the mechanism changes over time in response to a changing economic environment: David et al (2005) used Bayesian learning to optimize the rules of a single-sided auction mechanism in cases where agents are constrained to discrete bid prices, and Pardoe et al (2005) proposed a more general reinforcement-learning approach and used simulation to validate their framework in a setting in which the seller adapts their reserve price in order to optimise revenue when faced with an unknown population of loss-averse bidders.

In this paper we focus on evolutionary and iterative approaches to mechanism design, such as Conitzer and Sandholm (2007). The central theme of our work is that just as choice of strategy is not a static problem, since agents may be constrained in their adjustment of

strategy over time, neither is mechanism-design; mechanism designers may also be constrained in their choice of mechanism rules, for example there may be legacy infrastructure that prevents an institution such as a large stock exchange from radically altering its auction rules overnight. Just as constraints on strategy adjustment lead to *evolutionary* game theory, constraints on mechanism adjustment lead to *evolutionary* mechanism-design. We might think of the market institutions that we observe today as the equilibrium outcome of a co-evolutionary process not just between individual strategies, but a co-evolution between strategy and mechanism. Young (2001) was among the first to propose this idea with a view to *explaining* the evolution of economic institutions and social norms, and in our earlier work we used a similar idea to *design* economic institutions (Phelps et al, 2002a). In the next section we review this earlier work. Although this earlier approach has several shortcomings, the central ideas are instructive in understanding the rationale behind an iterative approach to mechanism design.

3.1 Co-evolution of auction mechanisms and trading strategies

In our earlier work (Phelps et al, 2002a) we approached mechanism design as an iterative process by simulating an co-evolutionary arms-race between two populations: on the one hand, a population of mechanisms whose fitness is determined by social welfare, and on the other hand a population of trading strategies whose fitness is determined by the utility of individual agents. The rationale behind this approach is that as strategy population evolved it would be able to find strategies that manipulate the mechanism in the favour of individual agents at the expense of social welfare. However, as non-truthful strategies become more prevalent, we might expect the mechanism population to counter-adapt by finding alterations to the mechanism rules that are able to restore social welfare. In these early co-evolutionary experiments, we hoped that over time incentive-compatible mechanism rules would evolve that were robust against a wide variety of trading strategies, in much the same way that prey populations adapt robust defenses against predator populations in co-evolutionary arms races in nature (Dawkins and Krebs, 1979; Valen, 1973).

This approach is similar in philosophy to that of Conitzer and Sandholm (2007):

“We start with a naïvely designed mechanism that is not strategy-proof (for example, the mechanism that would be optimal in the absence of strategic behavior), and we attempt to make it more strategy-proof. Specifically, the approach systematically identifies situations in which an agent has an incentive to manipulate, and corrects the mechanism to take away this incentive. This is done iteratively, and the mechanism may or may not become (completely) strategy-proof eventually.” (Conitzer and Sandholm, 2007, p. 2)

and reflects the iterative nature of real-world mechanism design:

“Real-world mechanisms are often initially naïve, leading to undesirable strategic behavior; once this is recognized, the mechanism is amended to disincent the undesirable behavior. For example, some naïvely designed mechanisms give bidders incentives to postpone submitting their bids until just before the event closes (i.e., sniping); often this is (partially) fixed by adding an activity rule, which prevents bidders that do not bid actively early from winning later.” (Conitzer and Sandholm, 2007, p. 2)

Indeed, the design of real-world market mechanisms usually proceeds incrementally, as accounts of the design of the rules for the US Federal Communications Commission's auctions of Personal Communications Services operating licences make clear (Nik-Khah, 2005; Plott, 1997).

In our earlier work we focussed on one specific aspect of the design of a double-auction market: the transaction pricing rule. The transaction pricing rule sets the price of any given transaction in the market as a function of the *bid* and *ask* offer prices submitted by buyers and sellers respectively. There are many different rules for setting transaction prices depending on the scenario at hand. For example, a discriminatory pricing rule sets transaction prices $P \subset 2^{\mathbb{R}}$:

$$P_i = MB_i \cdot k + MS_i \cdot (1 - k) \quad (2)$$

where $k \in [0, 1]$ is a constant chosen by the market designer, and the multi-sets MB and MS denote the prices of the matched⁶ bids and asks in the market respectively. A uniform pricing rule, on the other hand, sets the same transaction price $p \in \mathbb{R}$ for all matched offers:

$$p = eq_a \cdot k + eq_b \cdot (1 - k) \quad (3)$$

where:

$$\begin{aligned} eq_a &= \min(\min(MS'), \min(MB)) \\ eq_b &= \max(\max(MS), \max(MB')) \end{aligned}$$

and MB' and MS' denote the unmatched bid and ask prices respectively. Note that in case of a single-sided auction where we have a single seller with a reservation price, the $k = 1$ and $k = 0$ rules correspond to a 1st-price and 2nd-price (Vickrey) auction respectively (Wurman et al, 1998).

In a private-values auction scenario, offer prices can be thought of as signals⁷ from the traders expressing their valuation for the resource being traded. The difficulty the auctioneer faces in allocating the resource to those who value it most highly (i.e. achieving an optimal allocation or maximum market efficiency) is that these signals cannot necessarily be relied upon to be truthful; agents might misreport their valuations in order to make profit at the expense of others. The traditional approach to this problem is to design *incentive-compatible* mechanisms which have the property that the best strategy for every agent is to report their valuation truthfully.

In the case where only a single unit of commodity is traded per transaction, the pricing rules above can be shown to be efficient and incentive-compatible for sellers when $k = 1$, and similarly for buyers for $k = 0$. However, there is no value of k in a uniform-price mechanism for which the mechanism is simultaneously incentive-compatible and efficient

⁶ Offers are *matched* when the bid (buy) price meets or exceeds the ask (sell) price

⁷ The term "signal" in this context derives from the theory of *signalling games* (Spence, 1973). Although strictly speaking an auction is not a signalling game, the two are very strongly related; in a signalling game the agents move first and then the institution responds, whereas in a mechanism design scenario the institution offers a set of moves to agents who then respond (Dutta, 2001, p. 395). Thus although auctions are not strictly signalling games, it can still be intuitive to think in terms of signals; by forcing agents to back up their value claims with hard cash the mechanism designer can encourage *honest signalling*. Interestingly, signalling games have also been studied in evolutionary biology in the context of the *handicap principle* (Zahavi and Zahavi, 1997; Bullock, 1997). In the scenario under discussion, bids — that is, signals of valuation backed up with hard cash — can be thought of as "handicaps" which lead to honest signaling.

for all traders when multiple units can be traded (Wurman et al, 1998; Williams, 1988); i.e. there is no mechanism within this space which is strictly optimal in the theoretical sense. This result suggested to us the possibility of using heuristic search to find mechanisms that satisfied different design objectives.

In (Phelps et al, 2002a) a co-evolutionary algorithm (Hillis, 1992) was used to simulate an evolutionary “arms-race” between populations of trading strategies and a separate population of transaction pricing rules (the mechanism population). Individuals in each population were represented as Lisp expressions and evolved using Genetic Programming (Koza, 1993). The fitness function for the strategy populations was a function of the individual profits of traders playing those strategies, and the fitness function for the pricing rule population was a function of the overall market efficiency achieved by an auctioneer using that rule against the current strategy populations.

Despite some promising preliminary results, it was found that this approach suffered from a number of drawbacks. The main drawback was that the co-evolving system rapidly descended into suboptimal auction mechanisms if the mechanism population was not artificially seeded with individuals with a minimum-level of initial fitness. In cases where the mechanism population started from extremely low fitness individuals, such as pricing rules which set the transaction price at 0 regardless of the signals arriving from traders, the strategy populations would try and fit to these artificially low-fitness mechanisms and evolve to a state where their bids were meaningless. Meanwhile the mechanism population would be unable to discover more rational rules which worked with the existing “broken” trading strategies. Therefore the trading strategies could not evolve to work with more rational mechanisms and so on.

It is often instructive to analyse co-evolutionary processes in game-theoretic terms, since in a co-evolutionary interaction the fitness assigned to any given individual depends on the joint actions of the other individuals with which it interacts in a very similar manner to an evolutionary game⁸. When we co-evolve auction mechanisms and trading strategies we are implicitly defining a game⁹ between two players: the *mechanism player* on the one hand, and the *trader player* on the other. Each player attempts to maximise their payoff (analogous to maximising fitness); in our present scenario the mechanism player attempts to maximise market efficiency *EA*, whereas the trader player attempts to maximise their own utility. Note that if the selection function of our co-evolutionary algorithm picks individuals from each population based on a stochastic function of fitness rather than phenotype, then we are implicitly modelling a game of *imperfect information*. Although not every game possesses a dominant-strategy, we know that *all* games possess at least one Nash *equilibrium* in which the strategy adopted by every player is a best-response to every other player’s strategy. Consider a hypothetical equilibrium for our game at hand in which the mechanism population chooses a clearing-rule which sets the transaction price of all trades at a constant value which is independent of the trader’s shout price, and in response the trader player adopts a strategy of always submitting shouts with zero prices. Depending on the distribution of trader valuations, a rule which sets transaction prices close to the expected equilibrium price

⁸ Note that this applies regardless of whether we intuitively think of our original problem as a game. Game theory is simply a mathematical tool that allows us to study co-dependent optimization problems — that is, what potential solution should we choose given that our choice will influence the solution of other optimizers and vice versa. This is precisely the scenario instantiated by a co-evolutionary algorithm; hence game-theory is an invaluable theoretical tool in understanding the properties of co-evolutionary systems.

⁹ For conciseness and simplicity, in this section only we shall assume that many trading agents are under control of the single notional trader player, and that all the agents adopt the same strategy that is specified by the trader player at any given time.

would achieve a reasonable expected payoff $E(EA) \approx 1$ for the mechanism player. From an external mechanism designer's point of view this clearing rule is clearly brittle and undesirable, especially if the variance in valuations and hence in efficiency is large. However, this hypothetical situation would be very hard to leave once we arrive at it, since if the mechanism player attempts to switch to conventional clearing rules which set transaction prices as a function of shout prices, it will be faced with the issue that all shout prices are 0. Similarly, the trader player cannot improve their payoff by unilaterally switching to any other strategy since their payoff is no longer a function of their shout price. This situation is a game-theoretic *equilibrium* of the mechanism versus trader game.

If equilibria such as these have large basins of attraction¹⁰ under the dynamics of our co-evolutionary process, then we should not be surprised if our co-evolutionary algorithm converges on them. Indeed, this is the cause of one of the major problems that was encountered in the work reported in (Phelps et al, 2002a) — the co-evolutionary algorithm sometimes converged on what appeared to be game-theoretic equilibria¹¹, but it is not clear that the theoretical equilibrium solutions of the mechanism versus trader game are always desirable from a mechanism *design* perspective, as illustrated by the above example.

However, the Nash equilibria of the mechanism versus trader game are useful solutions to a different, but interesting, problem. If we are modelling a process in which multiple competing market institutions asynchronously adjust their rules over repeated interactions in response to observed trader strategies in the real world, and vice versa, (analogous to the scenario analysed by (Roth and Ockenfels, 2002) in which they compare two competing online auction formats: eBay and Amazon), then we might expect equilibrium solutions such as the fixed-price clearing rule to be the rational end result. It is not inconceivable, for example, that the reason that we continue to see a prevalence of fixed-price institutions such as bricks-and-mortar shops for selling consumer goods in real market places, despite the possibility of dynamically-priced institutions such as eBay¹², is due to fact that fixed-price institutions are an equilibrium solution of the real-life co-evolution between market mechanism and trader behavior. For example, consumers may be unable to switch from a fixed-price to an auction market for their required good since one may not exist yet, and correspondingly it may be very difficult for a startup to create an online auction market in the absence of existing traders on either side of the market. This is consistent with a view proposed by Philip Mirowski (Mirowski, 2001, pp.536–545) of economic marketplaces as complex ecologies. Some markets, such as garage sales, have relatively simple rules and procedures, while others, such as financial futures markets, are, by comparison, very complex. Yet all manage to co-exist, with each type of market, apparently, finding its own niche in which to survive and prosper. Indeed, the oldest markets have survived for hundreds of years without adopting the sophisticated rules evident in some newer markets. The behaviours of the participants in the different markets are, as one would expect, different. One challenge for computational economics, says Mirowski, is to explain this diversity of mechanisms, how it has arisen and how it is maintained.

As well as accounting for historical and present observations of actual market behavior, this analysis could also be *normative*; we might *recommend* that retailers adopt a fixed-price mechanism based on the fact that it is a best-response to the likely status quo. In this case we might interpret our solution as “the optimal” one in some sense.

¹⁰ See page 19.

¹¹ Interestingly the equilibrium state that we ended up with was close to the equilibrium computed in Zhan and Friedman (2005) where traders are restricted to constant mark-ups.

¹² <http://www.ebay.com/>

3.2 Mechanism Design as iterative heuristic optimisation

In the previous section we saw that co-evolutionary algorithms are natural models of games of imperfect information, or simultaneous move games. The previous experiments could be thought of as an analysis of evolutionary mechanism design in the case that the mechanism designer and the traders are simultaneously attempting to anticipate the choice of the other.

However, it is sometimes more natural to view real-world mechanism *design* as a *sequential* iterative process involving a single institution (Conitzer and Sandholm, 2007; Phelps, 2008). In this case, the considerations from the previous section do not apply, since the mechanism designer is given the opportunity to move first by announcing their mechanism rules publicly to the trader population, who then respond by placing shouts in the mechanism. In this scenario we no longer have a repeated simultaneous-move game, instead we have a 2-move extensive-form game. In the first move the mechanism player announces their mechanism rules with *perfect information*, and in the second move the trader player responds by placing shouts. In contrast to the previous section, in this scenario the trader player does not have to attempt to “anticipate” the move made by the mechanism player; rather it can form its strategy conditionally based on the mechanism rules chosen by the mechanism player. Thus, as a mechanism designer we should choose the optimal mechanism rules in the sense that the chosen rules optimise our design objectives when the trader player plays their best strategy *under that particular chosen mechanism*.

This scenario is not straight-forwardly modelled by a standard co-evolutionary algorithm; rather it is more natural to view it as a non-co-evolutionary optimisation problem in which we evaluate each potential mechanism by computing the values of our design objectives when traders play their best strategy for our candidate mechanism. An outline of such a process is summarised in the pseudocode below.

```

input : A set of initial heuristic strategies  $S$ , and a legacy mechanism  $\mu$ 
1 repeat
2    $\mathbf{x} \leftarrow$  frequency of each strategy observed in vivo;
3    $S \leftarrow S \cup \{ \text{strategies observed in vivo} \}$ ;
4    $\Lambda \leftarrow$  space of feasible variants of  $\mu$ ;
5    $\mu \leftarrow \arg \max_{\mu^* \in \Lambda} \text{EvaluateDesignObjectives}(\mu^*, S, \mathbf{x})$ ;
6   implement rules defined by  $\mu$ ;
7 until forever ;

```

We start with an initial set of auction rules comprising a real-life (*in-vivo*) mechanism μ , in which we observe a set of trading strategies S . We then update our analysis based on current observations of the real market (step 3); and finally we choose new mechanism rules that maximise our design objectives based on *in-silico* simulations of our proposed new design (step 5), before iterating the design cycle.

In (Phelps et al, 2003) we described how step 5 could be automated in order to evaluate market mechanisms with arbitrary design objectives and domain assumptions. Firstly we restricted attention to varying the parameter k in the transaction pricing rule given in equation 3, and evaluated a fitness function which combined two different objectives in a weighted sum: i) the extent to which the market structure favoured particular groups of traders (buyers or sellers); and (ii) the efficiency of the market. For each value of k , the objective function was assessed by running agent-based simulations with randomly-drawn valuations in order to estimate expected values of these metrics.

Figure 1 shows our empirical results for a double-auction obtained from simulations populated with sixty risk-neutral agents homogeneously equipped with a strategy based on the variant of the Roth-Erev reinforcement learning algorithm (Erev and Roth, 1998) described in (Nicolaisen et al, 2001). The Roth-Erev strategy (abbreviated hereon as RE) discretizes the space of possible offer prices, yielding a set of possible *actions* which can be thought of as pure strategies in a game of imperfect information. It then chooses actions stochastically with probabilities weighted in proportion to the cumulative past payoffs accrued to a particular action, in much the same way that fitness-proportionate selection operates in a co-evolutionary genetic algorithm. A potential advantage of the Roth-Erev algorithm over co-evolutionary algorithms is that the model has been tested against human data from controlled laboratory experiments¹³.

As can be seen from Figure—1, the optimal mechanism under this model is a pricing rule with $k = 0.5$. One of the restrictions of this approach is that we are constrained by our ability to parameterise the auction design space. Although much progress has been made towards this goal, both in the general case (Wurman et al, 2001) and in the specific case of the double auction (Niu et al, 2008), it is interesting to consider whether it is possible to view mechanism design as a less restricted search problem by using approaches which view mechanism design as an open-ended search for *symbolic* rules, thus allowing possible solutions which have not been a-priori categorized. In (Phelps et al, 2003) we extended the search space to all possible transaction pricing rules expressed as a symbolic linear function of bid, ask and quote prices. This space was searched heuristically using Genetic Programming Koza (1993), and our evolved symbolic rule was close to the $k = 0.5$ solution suggested by the more restrictive search, thus validating this as a potential approach to automated mechanism design.

¹³ Although this claim remains controversial since the model has many free parameters which need to be calibrated.

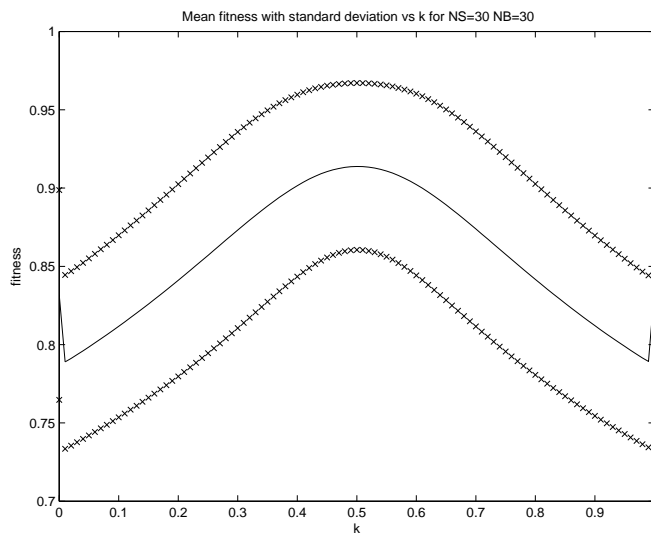


Fig. 1 Fitness F^i (with one standard deviation) plotted against k for a market with 60 traders.

In the case of single-sided auctions for a single unit of commodity, the revenue equivalence theorem (Krishna, 2002) states that if our design objective is restricted to maximising the revenue of the seller, then any value of k will suffice to achieve our design criterion provided that all bidders are risk-neutral and valuations are determined homogeneously by drawing signals from a random distribution that is common to all agents. However these assumptions are often violated in practice. Byde (2003) used a similar approach to ours in which he used a co-evolutionary genetic algorithm in place of reinforcement-learning to model the strategic behavior of agents in single-sided auctions. His goal was to search for values of k within the space of single-sided mechanisms, with a view to finding revenue-maximising auctions in scenarios where agent's risk attitudes and valuations are determined *heterogeneously*. By varying parameters governing risk preferences and commonality of valuations he was able to demonstrate that values of k intermediate between $k = 0$ and $k = 1$ are optimal depending on our assumptions about the specifics of the economic environment to hand.

3.2.1 Empirical game-theory in mechanism design

There are several drawbacks with the type of approach discussed in the previous section, since we either: i) assume that the particular strategy (eg Roth-Erev) used in the agent-based model constitutes a strategic equilibrium which other strategies are unable to invade; or ii) if we use co-evolutionary search to explore a larger space of strategies, we assume that the asymptotic states of the co-evolutionary algorithm similarly constitute Nash equilibria. As discussed in Section 2.1.3, neither of these assumptions are guaranteed to hold.

To deal with this criticism, in later work we used empirical game-theory¹⁴ to analyse the effect of changing from a uniform to a discriminatory pricing rule. Within this context, the analysis from the previous section can be understood as an empirical game-theoretic analysis restricted to a *single* heuristic strategy: the Roth-Erev strategy. However, in (Phelps et al, 2006b) we extended our analysis to four classes of strategy: i) the Roth-Erev strategy (RE) described on p. 17, ii) the truth-telling strategy (TT) which bids at the agent's valuation, iii) the Gjerstad-Dickhaut strategy (GD) (Gjerstad and Dickhaut, 2001) and iv) Kaplan's sniping strategy (TK) (Friedman and Rust, 1991). As in Walsh et al (2002) we used the replicator-dynamics differential equation to model how agents switch between these strategies in response to observed payoffs:

$$\dot{m}_j = [u(e_j, \mathbf{m}) - u(\mathbf{m}, \mathbf{m})] m_j \quad (4)$$

where \mathbf{m} is a mixed-strategy vector, $u(\mathbf{m}, \mathbf{m})$ is the mean payoff when all players play \mathbf{m} , and $u(e_j, \mathbf{m})$ is the average payoff to pure strategy j when all players play \mathbf{m} , and \dot{m}_j is the first derivative of m_j with respect to time. Strategies that gain above-average payoff become more likely to be played, and this equation models a simple co-evolutionary process of mimicry learning, in which agents switch to strategies that appear to be more successful.

For any initial mixed-strategy we can find the eventual outcome from this co-evolutionary process by solving $\dot{m}_j = 0$ for all j to find the final mixed-strategy of the converged population. This model has the attractive properties that: (i) all Nash equilibria of the game are stationary points under the replicator dynamics; and (ii) all Lyapunov stable states (Lyapunov, 1966) and interior limit states are also Nash equilibria (Weibull, 1997, pp. 88–89).

Thus the Nash equilibrium solutions are embedded in the stationary points of the direction field of the dynamics specified by equation 4. Although not all stationary points are

¹⁴ As pioneered by (Walsh et al, 2002); see Section 2.2 for a full description and literature review.

Nash equilibria, by overlaying a dynamic model of learning on the equilibria we can see which solutions are more likely to be discovered by *boundedly-rational* agents. Those Nash equilibria that are stationary points at which a larger range of initial states will end up, are equilibria that are more likely to be reached (assuming an initial distribution that is uniform and that the replicator dynamics are an accurate description of the switching process).

For a subset $n = 3$ of our strategies we can geometrically project the space of all possible mixed-strategy vectors, called the unit-simplex¹⁵, onto a two dimensional triangle denoted Δ^2 whose vertices correspond to the pure strategies $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 1)$. This gives us an intuitive visualisation of the dynamics of the learning process; by plotting the time-evolution of equation 4 we can identify the switching between our heuristic strategies. Figure 2 shows the direction field when we consider evolutionary switching between the three strategies TT, RE and GD, in a $k = 0.5$ uniform-price clearing house market populated by twelve agents which are selected at random from a larger population of traders on each play of the game.

The direction field gives us a map which shows the trajectories of strategies of learning agents engaged in repeated interactions, from a random starting position. Thus, for Figure 2, each agent participant has a starting choice of 3 pure strategies (TT, RE and GD) and any mixed (probabilistic) combination of these three. The arrows indicate the direction of convergence when any such strategy is adopted. The three pure strategies (here, TT, RE and GD) are represented by the three vertexes of the simplex. A point on an external edge of the simplex represents a mixed strategy comprising two of the three pure strategies, and a point strictly inside the simplex represents a mixed strategy comprised of all three pure strategies. Thus, for example, the point on the left-most edge between the vertexes labeled TT and RE which is one-third the way from the vertex labeled TT represents a mixed strategy where strategy TT is chosen 66.7% of the time, strategy RE is chosen 33.3% of the time, and strategy GD not chosen at all; this position on the simplex is denoted $(66.7, 33.3, 0)$. A vector (a line with an arrow) shows the likely direction of strategic play from any given initial position. In other words, if the arrows converge on some point in the simplex, this strategy represented by that point is the end-point of repeated interactions as the game proceeds.

The Nash equilibrium solutions of the heuristic-game are embedded in the stationary points of the direction field of the dynamics specified by equation 4. In Figure 2 there are two equilibria: a pure-strategy GD equilibrium at the bottom right of the simplex, and a mixed-strategy comprising a high probability of TT and a low probability of RE near the bottom left of the simplex. Although not all stationary points are Nash equilibria, by overlaying a dynamic model of learning on the equilibria we can see which solutions are more likely to be discovered by *boundedly-rational* agents. Those Nash equilibria that are stationary points at which a larger range of initial states will end up, are equilibria that are more likely to be reached (assuming that initial values of \mathbf{m} are uniformly distributed over Δ , and that the replicator dynamics is an adequate model of the switching process); in the terminology of dynamic systems they have a larger *basin of attraction*. The basin of attraction for a stationary point is proportion of mixed strategies in Δ which have flows terminating at that point¹⁶. The larger the basin, the larger the region of strategy-space which leads to the attractor, and the more *attainable* the corresponding equilibrium (Bullock, 1997), and hence the higher its probability of occurring.

¹⁵ $\Delta^{n-1} = \{\mathbf{x} \in \mathbb{R}^n : \sum_i^n x_i = 1\}$

¹⁶ Intuitively, it can be helpful to conceptualise basin size as the *volume* of the state space which terminates at the attractor. However, strictly speaking this definition is not accurate. For example, if we have chaotic dynamics then a strange attractor may capture many flows, but the *volume* of its basin will be zero.




Fig. 2 3-dimensional replicator dynamics direction field for a 12-agent clearing-house auction with the three strategies RE, TT and GD.

This analysis is highly valuable for the purposes of mechanism design, since different equilibria will yield different outcomes and different values of our design objectives, such as market efficiency, and we would like to weight these according to their likelihood. In other words, we would like to compute the size of the basin of attraction of each equilibrium, in order to arrive at a probability of the equilibria actually occurring, and use this to calculate the expected value of our design metrics. We can then weight the design objectives for each mechanism according to the probability distribution over equilibria, which will allow us to provide more realistic estimates for our design metrics.

In (Phelps et al, 2006b) we used numerical methods to estimate the basin size of equilibria for $n = 4$ strategies for two different variants of the double-auction: i) a clearing-house (CH) in which offers from all agents are collected before the market is cleared; and ii) a continuous double-auction (CDA) in which the market is cleared as new offers arrive. Our main results are reproduced in Table 1.

These results give us probabilities over outcomes, and hence the ability to assess the design of each mechanism. The value of EA in each cell of Table 1 shows our estimate of the expected efficiency¹⁷ of the mechanism. This is computed by weighting the pure-strategy payoffs U according to the probability of the pure strategy being played. For example, in the case of CDA with 6 agents, we see that there are two possible equilibria. The first equilibrium, $(0, 0, 0.84, 0.16)$, has a probability of 0.97 of being adopted. In this equilibrium the

¹⁷ See Phelps (2008) for a formal definition of market efficiency (EA).

CH		
4 agents	6 agents	12 agents
$\beta(0, 0, 1, 0) = 0.39$ $\beta(0, 0, 0, 1) = 0.61$	$\beta(0, 0, 1, 0) = 0.31$ $\beta(0, 0, 0, 1) = 0.69$	$\beta(0, 0, 0, 1) = 1$
$U = (1.00, 0.90, 1.00, 1.00)$ $EA = 1.00$	$U = (1.00, 0.92, 1.00, 1.00)$ $EA = 1.00$	$U = (1.00, 0.93, 1.00, 1.00)$ $EA = 1.00$
CDA		
4 agents	6 agents	12 agents
$\beta(0, 0, 0, 1) = 1$	$\beta(0, 0, 0.84, 0.16) = 0.97$ $\beta(0, 0, 0, 1) = 0.03$	$\beta(0, 0, 0.8, 0.2) = 1$
$U = (0.89, 0.86, 0.98, 0.89)$ $EA = 0.89$	$U = (0.85, 0.88, 0.98, 0.86)$ $EA = 0.96$	$U = (0.85, 0.89, 0.99, 0.90)$ $EA = 0.97$

Table 1 Heuristic-strategy equilibria over (TT, RE, GD, TK) for CH versus CDA

strategy GD has a probability 0.84 whereas the strategy TK has a probability of 0.16. By examining the payoffs to each of these strategies we can compute the expected efficiency of the mechanism in this equilibrium: $0.84 \times 0.98 + 0.16 \times 0.86 = 0.96$. In the second equilibrium we see that the strategy TK has a probability 1 of being played, hence the efficiency of this second equilibrium is 0.86. We then weight our overall efficiency according to the probability of each equilibrium: $0.96 \times 0.97 + 0.86 \times 0.03 = 0.96$.

For the most part efficiency outcomes are deterministic — there is either a unique equilibrium that captures the entire simplex or all equilibria yield the same efficiency. The exception is the CDA with 6 agents. Here we have a mixed TK and GD equilibrium with efficiency $EA = 0.97$ versus a pure TK equilibrium with a *significantly lower* efficiency of $EA = 0.86$. Since the TK equilibrium has a very small basin of attraction $\beta = 0.03$ we conclude that the lower efficiency outcome is not very likely, and hence if we have no prior knowledge of existing strategy frequencies in the trading population at large we assume a uniform distribution over starting points in the simplex and conclude that our efficiency is still likely to be very high. However, in the case where we do have prior knowledge about the frequency of strategies, e.g. we are tasked with evaluating a proposed choice of a continuous-clearing rule for a six-agent marketplace in which we *already* observe high proportion of sniping, then we might conclude that the pure TK equilibrium is much more likely to be reached (since we will be starting within its attractor), and thus we might recommend that CH clearing is used instead in order to avoid the probable efficiency hit predicted by our analysis. This hypothetical design tweak would yield an efficiency gain of $0.97 - 0.86$, or 11 percentage points, at the expense of transaction throughput. Thus by analysing the strategic *dynamics* of a proposed mechanism, we can perform *evolutionary* mechanism design whereby we make design decisions under *legacy* constraints (in this hypothetical scenario our legacy constraint is an existing marketplace populated by snipers). Evolutionary mechanism design is analogous to evolutionary game theory in that just as players may be constrained to gradually adjust their strategies, similarly mechanisms cannot always make instantaneous adjustments in their rules irrespective of what strategies are currently in play.

3.3 Iterative refinement of heuristic-strategy analysis

One of the criticisms of empirical game-theoretic analysis used in the previous section is that it is highly sensitive to the set of heuristic strategies S , which can never be truly comprehensive for an initial design. In recent work (Phelps et al, 2008, 2006a) we attack this problem by introducing an algorithm called FISH+ for searching for new heuristic strategies

within an iterative mechanism-design context. Conventional approaches to strategy optimisation use a strategy’s payoff as the basis of an objective function. In contrast, the objective function in our approach is the strategy’s likelihood of being adopted in equilibrium play, as estimated from basin size. In (Phelps et al, 2008) we argue that the latter approach is more suited to heuristic-search of the strategy space in economic scenarios.

The basic idea is to use a sensitivity analysis to identify which of our existing strategies could obtain a larger basin size, if its expected payoffs could be increased slightly. We then use heuristic search to find variants of this strategy which yield high “market-share”: that is their probability of occurring in equilibria, given by the equilibria basin size weighted by the frequency of adoption specified by the equilibrium mixed-strategy vector. This is formalised below:

$$F(i, S, [H]) = \sum_{\mathbf{x} \in \epsilon_{[H]S}} \beta_{[H]}(\mathbf{x}, M) \cdot x_i \quad (5)$$

where: i is the index of the candidate heuristic strategy being evaluated from amongst the set of heuristic strategies S with heuristic payoffs $[H]$, $\beta_{[H]}$ denotes the basin size of an equilibrium in the game defined by payoffs $[H]$, and $\epsilon_{[H]S}$ is the set of heuristic equilibria¹⁸:

$$\epsilon_{[H]S} = \{\mathbf{x} \in \Delta^{|S|} : \beta_{[H]}(\mathbf{x}, M) > 2 \times 10^{-2}\}$$

By applying this search iteratively, we can refine our initial heuristic-strategy analysis of the marketplace and find mixtures of strategies that yield large basin sizes which are stable to payoff perturbation:

```

input : A set of heuristic strategies  $S = \{s_1, s_2, \dots, s_n\}$  for some mechanism  $\mu$ 
output: A refined set of heuristic-strategies
1  $[H] \leftarrow \text{GetHeuristicPayoffMatrix}(S, \mu)$ ;
2 repeat
3    $\hat{F} \leftarrow \max_{i=1 \dots n} F(i, S, [H])$ ;
4   for  $i \leftarrow 1$  to  $n$  do
5      $[H]' \leftarrow \text{perturb payoffs in } [H] \text{ in favour of } s_i$ ;
6     if  $F(i, S, [H]') > \hat{F}$  then
7        $\hat{F} \leftarrow F(i, S, [H]')$ ;
8        $i^* \leftarrow i$ ;
9        $\hat{OS} \leftarrow s_i$ ;
10    end
11  end
12  if  $\hat{F} < F(i^*, S, [H])$  then return  $S$ ;
13   $\Pi \leftarrow \text{create a search space based on generalisations of } \hat{OS}$ ;
14   $OS \leftarrow \arg \max_{s^* \in \Pi} F(1, s^* \cup S, \text{GetHeuristicPayoffMatrix}(s^* \cup S, \mu))$ ;
15   $S \leftarrow OS \cup S$ ;
16   $[H] \leftarrow \text{GetHeuristicPayoffMatrix}(S, \mu)$ ;
17   $S \leftarrow \text{eliminate dominated strategies from } S \text{ based on } [H]$ ;
18 until forever ;

```

Algorithm 1: FiSH+

¹⁸ Here we use the heuristic that stationary points with a basin size above 2% are asymptotically stable and hence equilibria

To a mechanism designer, this latter state of affairs is particularly attractive, *since larger, more stable basin sizes correspond to more deterministic, and hence predictable behaviour*. In a legacy mechanism design scenario, if we are able to provide an equilibrium analysis over existing strategies which demonstrates similarly clear-cut equilibria, then we may be able to convince participants that these are the best-response strategies that their competitors are likely to adopt, and therefore that they should adopt also. If we then make the algorithms corresponding to our *new* heuristic strategies freely available to participants, and if they believe our equilibrium analysis, then they are likely to play our prescribed strategies, thus bringing about our predictions, and hence maximising our design objectives. By finding new strategies with large stable attractors, we make our equilibrium analysis more believable to participants. This is analogous to incentive-compatibility in a conventional mechanism design scenario, where it is clear to participants that TT is the traders' best-response to the mechanism: in an incentive-compatible mechanism TT is a "freely-available" strategy with a large attractor. In realistically complex mechanisms such as the double-auction, TT is dominated. However by applying the FISH+ algorithm we can find analogs of TT for complex mechanisms.

Of course, in our new equilibria, our existing mechanism rules may no longer maximise our design objectives. In this paper, we have described real-life mechanism design as an iterative process, and that is exactly how evolutionary mechanism design addresses this issue. Thus our revised pseudo-code for evolutionary mechanism design is as follows:

<pre> input : A set of initial heuristic strategies S, and a legacy mechanism μ 1 repeat 2 $S \leftarrow \text{FISH+}(S, \mu)$; 3 publicise S to participants; 4 $\mathbf{x} \leftarrow$ frequency of each strategy observed in vivo; 5 $S \leftarrow S \cup \{ \text{strategies observed in vivo} \}$; 6 $\Lambda \leftarrow$ space of feasible variants of μ; 7 $\mu \leftarrow \arg \max_{\mu^* \in \Lambda} \text{EvaluateDesignObjectives}(\mu^*, S, \mathbf{x})$; 8 implement rules defined by μ; 9 until forever ; </pre>
--

Algorithm 2: Evolutionary mechanism design

4 Summary

Economists have long used idealized models of agent behaviour in order to understand market behaviour. AI practitioners have had to adapt these models in order to build actual agents, and the resulting engineering approach to agents' behaviour requires more sophisticated and complex models. Similarly, it has recently been understood that the idealized notion of a "free" market is not always applicable, since actual markets entail many rules that govern their operation. Building real markets entails an engineering approach just as does the building of real agents.

In this paper we have reviewed iterative methodologies for the engineering of market mechanisms, which we categorize under the banner of *evolutionary mechanism design*. This differs from traditional mechanism design, which is a static analysis based on rigidly defined design objectives, in which a theoretically pristine mechanism is launched into the world and

then remains forever in Nash equilibrium stasis. Evolutionary mechanism design, in contrast, attempts to take an *engineering* approach, and approach that is empirical, incremental, and partly automated. It is not theoretically beautiful, but it is able to take into account real-life ugliness: arbitrary multiple design objectives, *dynamic* adjustment to equilibrium, and constant feedback from an *in vivo* mechanism.

Just as with engineering methods for other complex real-world domains, such as software engineering, our initial analysis cannot be relied upon to be completely accurate and future-proof (Bittner and Spence, 2006; Beck, 1999). Therefore we continually update our analysis in response to feedback from the *in vivo* mechanism: in steps 4 and 5 we compare our predictions with actuality, and update our set of heuristic strategies S and their observed frequency in the population x .

Finally, the resulting status quo may not be optimal for our purposes; for example, we may be able to improve the likelihood of achieving certain design objectives, such as market efficiency or liquidity (transaction throughput) by making small adjustments in a subset of the space of mechanism rules, for example by adjusting parameters such as k in the market clearing rules (equations 2 and 3), as we saw in Section 3.2.

4.1 Future work

In this paper we have concentrated on the purely computational aspects of the method: that is, *in silico* analysis. In so doing, we have glossed over some of the challenges presented by the *in vivo* analysis of real-life market places, which may be considerable. For example, in Section 3.2.1, we saw how our design objectives were affected when we considered different weightings over the frequency with which sniping strategies were observed in the existing mechanism. In the case of a strategy such as sniping, it is relatively straightforward to determine which traders are adopting this strategy, provided that one has access to sufficient historical market data, since we can simply look at the timing of agents' shouts; Roth and Ockenfels (2002) provide just such an analysis of the eBay marketplace, which validates that steps 4 and 5 can be performed *in vivo* in the case of a single class of strategy.

However, inferring the existence of other classes of strategies in a real market presents a significant challenge, not least because the true valuation of each agent is not directly observable. Without any prior knowledge of an agent's valuation, it is very difficult to infer whether they are using a strategy even as simple as truth-telling. That is not to say, however, that making inferences about valuations is impossible, especially from the privileged vantage point of the agent controlling the mechanism, who potentially has full access to the history of traders' interactions with the market. We may, for example, be able to infer bounds on an agent's valuation by analysing the order statistics of their trade prices over small time periods; or by analysing their trading behaviour in alternative markets for the same commodity; or, in the case of an ascending auction format such as eBay, by observing the price at which runner-up bidders drop out of the auction. With estimates of valuations in hand, it would be possible in many cases to infer an agent's strategy. The reverse-engineering of valuations and strategies from market data is a promising area of research, both for those seeking to make profit, as well as for economists seeking to understand the dynamics of real-world marketplaces, and there is an emerging literature in this area (Engle-Warnick and Ruffle, 2001; Engle-Warnick, 2003) to draw upon.

Although it might be impractical in the context of an academic research programme to apply these *in vivo* methods in the context of a market such as a stock exchange, it may be possible to apply them to a markets such as the University of Iowa prediction markets

(Surowiecki, 2004). Prediction markets are exchanges with unique design considerations (Wolfers and Zitzewitz, 2004), and an interesting possible research programme would be to conduct a full *in vivo* case study of the application of evolutionary mechanism design to a real-life prediction market through several iterations of the design cycle.

References

- Anderson R (2001) Why information security is hard: An economic perspective. In: Proceedings of the 17th Annual Computer Security Applications Conference, IEEE Computer Society, Los Alimos, CA, USA, p 358
- Andrews M, Prager R (1994) Genetic programming for the acquisition of double auction market strategies. In: Kinnear, Jr KE (ed) *Advances in Genetic Programming*, MIT Press, chap 16, pp 355–368
- Angeline PJ, Pollack JB (1993) Competitive environments evolve better solutions for complex tasks. In: Forrest S (ed) *Genetic Algorithms: Proceedings of the Fifth International Conference (GA93)*
- Axelrod R (1997) *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*. Princeton University Press
- Beck K (1999) *Extreme Programming Explained: Embrace Change*. Addison-Wesley
- Bendor J, Swistak P (1997) The evolutionary stability of cooperation. *American Political Science Review* 91(2):290–307
- Bittner K, Spence I (2006) *Managing Iterative Software Development Projects*. Addison-Wesley
- Boutilier C, Shoham Y, Wellman MP (1997) Editorial: Economic principles of multiagent systems. *Artificial Intelligence* 94:1–6
- Bryant WDA (2000) Information, adjustment and the stability of equilibrium. Tech. Rep. 6/2000, Department of Economics, Macquarie University, Sydney, Australia
- Bullock S (1997) Evolutionary simulation models: On their character, and application to problems concerning the evolution of natural signalling systems. PhD thesis, University of Sussex
- Byde A (2003) Applying evolutionary game theory to auction mechanism design. In: Proceedings of the ACM Conference on Electronic Commerce 2003, pp 192–193
- Chatterjee K, Samuelson W (1983) Bargaining under incomplete information. *Operations Research* 31(5):835–851
- Clarkson GPE, Simon HA (1960) Simulation of individual and group behavior. *American Economic Review* 50:920–932
- Cliff D (1998) Evolving parameter sets for adaptive trading agents in continuous double-auction markets. In: *Agents-98 Workshop on Artificial Societies and Computational Markets*, Minneapolis, MN., pp 38–47
- Cliff D (2001) Evolution of market mechanism through a continuous space of auction-types. Tech. Rep. HPL-2001-326, HP Labs
- Cliff D, Bruten J (1997) Minimal-intelligence agents for bargaining behaviors in market-based environments. Tech. Rep. HPL-07-91, HP Labs, Bristol
- Conitzer V, Sandholm T (2002) Complexity of mechanism design. In: Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence (UAI-02), pp 103–110
- Conitzer V, Sandholm T (2007) Incremental mechanism design. In: Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)

- Dash R, Parkes D, Jennings N (2003) Computational mechanism design: A call to arms. *IEEE Intelligent Systems* 18(6):40–47, URL citeseer.ist.psu.edu/dash03computational.html
- d'Aspremont C, Gérard-Varet L (1979) Incentives and incomplete information. *Journal of Public Economics* 11:25–45
- David E, Rogers A, Schiff J, Karus S, Jennings NR (2005) Optimal design of english auctions with discrete bid levels. In: *Proceedings of the 6th ACM Conference on Electronic Commerce (EC'05)*, Vancouver, Canada, pp 98–107
- Dawid H (1999) The convergence of genetic learning in a double auction market. *Journal of Economic Dynamics and Control* 23:1545–1567, URL citeseer.nj.nec.com/438397.html
- Dawkins R, Krebs JR (1979) Arms races between and within species. *Proceedings of the Royal Society of London B* 205:489–511
- Dutta PK (2001) *Strategies and Games: Theory and Practice*. MIT Press
- Engle-Warnick J (2003) Inferring strategies from observed actions: a nonparametric, binary tree classification approach. *Journal of Economic Dynamics and Control* 27:2151–2170
- Engle-Warnick J, Ruffle B (2001) Inferring buyer strategies and their impact on monopolist pricing. Tech. Rep. 2001-W28, Economics Group, University of Oxford, Nuffield College
- Erev I, Roth AE (1998) Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review* 88(4):848–881
- Ficici SG (2004) *Solution concepts in coevolutionary algorithms*. PhD thesis, Brandeis University
- Ficici SG, Pollack JB (1998) Challenges in coevolutionary learning: Arms-race dynamics, open-endedness, and mediocre stable states. In: *Proceedings of ALIFE-6*
- Ficici SG, Pollack JB (2000) A game-theoretic approach to the simple coevolutionary algorithm. In: Marc Schoenauer, Kalyanmoy Deb, Günter Rudolph, Xin Yao, Evelyne Lutton, Juan Julian Merelo HPS (ed) *Parallel Problem Solving from Nature — PPSN VI 6th International Conference*, Springer Verlag, Paris, France
- Ficici SG, Pollack JB (2003) A game-theoretic memory mechanism for coevolution. In: et al ECP (ed) *LNCS 2723*, Springer-Verlag, pp 286–297
- Fisher FM (1983) *Disequilibrium Foundations of Equilibrium Economics*. Cambridge University Press
- Flood MM (1952) Some experimental games. Tech. Rep. RM-789, RAND Corporation, Santa Monica, CA
- Friedman D, Rust J (eds) (1991) *The Double Auction Market: Institutions, Theories, and Evidence* (Proceedings of the workshop on double auction markets held June 1991 in Santa Fe, New Mexico). Westview
- Fudenberg D, Mobius MM, Szeidl A (2003) Existence of equilibrium in large double auctions. Working paper, Department of Economics, Harvard University, Littauer Center, Cambridge, MA 02138
- Gintis H (2000) *Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Interaction*. Princeton University Press
- Gjerstad S, Dickhaut J (2001) Price formation in double auctions. *Lecture Notes in Computer Science* 2033:106
- Gode DK, Sunder S (1993) Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy* 101(1):119–137
- Goeree JK, Holt CA (2001) Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review* 91(5):1492–1422

- Greenwald AR, Stone P (2001) Autonomous bidding agents in the trading agent competition. *IEEE Internet Computing* 5(2):52
- Harsanyi JC (1967) Games with incomplete information played by Bayesian players: Part i. *Management Science* pp 159–182
- Hillis WD (1992) Co-evolving parasites improve simulated evolution as an optimization procedure. In: et al L (ed) *Proceedings of ALIFE-2*, Addison Wesley, pp 313–324
- Hsu W, Soo V (2001) Market performance of adaptive trading agents in synchronous double auctions. In: *Lectures Notes in Computer Science*, vol 2132, Springer Berlin, pp 108–121
- Huang P, Scheller-Wolf A, Sycara K (2002) A strategy-proof multiunit double auction mechanism. In: *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Bologna, Italy, pp 166–167
- Jackson MO (2003) Mechanism theory. In: Devigs U (ed) *Optimization and Operations Research*, The Encyclopedia of Life Support Science, EOLSS Publishers, Oxford, UK
- Jackson MO, Swinkels JM (2001) Existence of equilibrium in single and double private value auctions. Mimeo, California Institute of Technology
- Jackson MO, Swinkels JM (2005) Existence of equilibrium in single and double private value auctions. *Econometrica* 73(1):93–139
- Jordan P, Vorobeychik Y, Wellman MP (2008) Searching for approximate equilibria in empirical games. In: Padgham, Parkes, Müller, Parsons (eds) *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, Estoril, Portugal
- Jordan PR, Kiekintveld C, Wellman MP (2007) Empirical game-theoretic analysis of the TAC supply chain game. In: *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems*, Honolulu, Hawaii, pp 1188–1195
- Kadan O (2004) Equilibrium in the two player, k -double auction with affiliated private values. Tech. Rep. 12.2004, Nota di Lavoro
- Kagel JH, Roth AE (eds) (1995) *The Handbook of Experimental Economics*. Princeton University Press
- Klemperer P (2002) How (not) to run auctions: the European 3G telecom auctions. *European Economic Review* 46:829–845
- Klemperer P (2004a) *Auctions: Theory and Practice*. Princeton University Press
- Klemperer P (2004b) Why every economist should learn some auction theory. In: *Auctions: Theory and Practice*, Princeton University Press
- Koza JR (1993) *Genetic Programming: on the programming of computers by means of natural selection*. MIT Press
- Krishna V (2002) *Auction Theory*. Harcourt Publishers Ltd.
- Lyapunov AM (1966) *Stability of motion*. Academic Press, New York and London
- MacKenzie DA (2003) An equation and its worlds: *Bricolage*, exemplars, disunity and performativity in financial economics. *Social Studies of Science* 33:831–868
- Maynard-Smith J (1982) *Evolution and the Theory of Games*. Cambridge University Press
- McAfee RP (1992) A dominant strategy double auction. *Journal of Economic Theory* 56:434–450
- Medio A, Gallo G (1992) *Chaotic Dynamics: Theory and Applications to Economics*. Cambridge University Press
- Mirowski P (2001) *Machine Dreams: Economics Becomes a Cyborg Science*. Cambridge University Press
- Myerson RB, Satterthwaite MA (1983) Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 28:265–281
- Nash J (1950) Equilibrium points in n -person games. In: *Proceedings of the National Academy of Sciences*, vol 36, pp 48–49

- Nicolaisen J, Petrov V, Tesfatsion L (2001) Market power and efficiency in a computational electricity market with discriminatory double-auction pricing. *IEEE Transactions on Evolutionary Computation* 5(5):504–523
- Nik-Khah E (2005) Designs on the mechanism: Economics and the FCC Auctions. Ph.d., University of Notre Dame, Notre Dame, Indiana, USA
- Niu J, Cai K, Gerding E, McBurney P, Parsons S (2008) Characterizing effective auction mechanisms: Insights from the 2007 TAC Mechanism Design Competition. In: Padgham, Parkes, Müller, Parsons (eds) *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, Estoril, Portugal, pp 1079–1086
- Noble J, Watson RA (2001) Pareto coevolution: Using performance against coevolved opponents in a game as dimensions for pareto selection. In: *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO-2001*, Morgan Kaufman, San Francisco, California, pp 493–50
- Papadimitriou CH (2001) Algorithms, games, and the internet. In: *Proceedings of the 33rd Symposium on Theory of Computing*, ACM Press, pp 749–753
- Pardoe D, Stone P, Saar-Tsechansky M, Tomak K (2005) Adaptive auctions: Learning to adjust to bidders. In: *Proceedings of the Fifteenth Annual Workshop on Information Technologies and Systems*
- Phelps S (2008) Evolutionary mechanism design. PhD thesis, University of Liverpool, Liverpool, UK., submitted July 2007.
- Phelps S, Parsons S, McBurney P, Sklar E (2002a) Co-evolution of auction mechanisms and trading strategies: Towards a novel approach to microeconomic design. In: *Proceedings of the Bird of a Feather Workshops, Genetic and Evolutionary Computation Conference, AAAI*, New York, pp 65–72
- Phelps S, Parsons S, McBurney P, Sklar E (2002b) Co-evolutionary mechanism design: A preliminary report. In: Padget J, Shehory O, Parkes D, Sadeh N, Walsh WE (eds) *Agent-Mediated Electronic Commerce IV: Designing Mechanisms and Systems*, Springer Verlag, pp 123–143
- Phelps S, Parsons S, Sklar E, McBurney P (2003) Using genetic programming to optimise pricing rules for a double auction market. In: *Proceedings of the workshop on Agents for Electronic Commerce*, Pittsburgh, PA
- Phelps S, Marcinkiewicz M, Parsons S, McBurney P (2006a) A novel method for automatic strategy acquisition in n-player non-zero-sum games. In: Nakashima H, Wellman MP, Weiss G, Stone P (eds) *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006)*, ACM, Hajodate, Japan, pp 705–712
- Phelps S, Parsons S, McBurney P (2006b) An evolutionary game-theoretic comparison of two double-auction market designs. In: Faratin P, Rodriguez-Aguilar JA (eds) *Agent-Mediated Electronic Commerce VI*, Springer Verlag, pp 101–114
- Phelps S, Parsons S, McBurney P (2008) A novel method for strategy acquisition and its application to a double-auction market. *IEEE Transactions on Systems, Man, and Cybernetics: Part B* *In submission*
- Plott CR (1997) Laboratory experimental testbeds: application to the PCS auction. *Journal of Economics and Management Strategy* 6(3):605–638
- Pollack JB, Blair AD (1998) Co-evolution in the successful learning of backgammon strategy. *Machine Learning* 32:225–240
- Preist C, van Tol M (2003) Adaptive agents in a persistent shout double auction market. Tech. Rep. HPL-2003-242, HP Laboratories Bristol

- Price TC (1997) Using co-evolutionary programming to simulate strategic behaviour in markets. *The Journal of Evolutionary Economics* 7:219–254
- Reeves DM, MacKie-Mason JK, Wellman MP, Osepayshvili A (2005) Exploring bidding strategies for market-based scheduling. *Decision Support Systems*
- Reny PJ, Perry M (2003) Toward a strategic foundation for rational expectations equilibrium. Working paper, University of Chicago
- Rosenschein JS, Zlotkin G (1994) *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT press
- Roth AE (2002) The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica* 70(4):1341–1378
- Roth AE, Ockenfels A (2002) Last-minute bidding and the rules for ending second-price auctions: Evidence from ebay and amazon auctions on the internet. *American Economic Review* 92(4):1093–1103
- Russell S, Norvig P (2003) *Artificial Intelligence: A Modern Approach*, 2nd edn. Prentice Hall
- Rustichini A, Satterthwaite MA, Williams SR (1994) Convergence to efficiency in a simple market with incomplete information. *Econometrica* 62(5):1041–1063
- Sandholm T (2003) Automated mechanism design: A new application area for search algorithms. In: *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*
- Satterthwaite M, Williams S (1993) The Bayesian theory of the k-double auction. In: Friedman, Rust (eds) *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Studies in the Sciences of Complexity, Westview, chap 4, pp 199–125
- Smith VL (1962) An experimental study of competitive market behavior. *Journal of Political Economy* 70:111–337
- Spence M (1973) Job market signaling. *The Quarterly Journal of Economics* 87(3):355–374
- Surowiecki J (2004) *The wisdom of crowds: why the many are smarter than the few and how collective wisdom shapes business, economies, societies, and nations*. Doubleday
- Tesauro G, Bredin JL (2002) Strategic sequential bidding in auctions using dynamic programming. In: *Proceedings of the first international joint conference on autonomous agents and multi-agent systems*, ACM, pp 591–598
- Tesauro G, Das R (2001) High-performance bidding agents for the continuous double auction. In: *Proceedings of the third ACM Conference on Electronic Commerce*, Tampa, Florida, USA, pp 206–209
- Tesfatsion L (2002) Agent-based computational economics: growing economies from the bottom up. *Artificial Life* 8(1):55–82
- Valen LV (1973) A new evolutionary law. *Evolutionary Theory* 1:1–30
- Varian HR (1995) Economic mechanism design for computerized agents. In: *Proceedings of the First USENIX Workshop on Electronic Commerce*, pp 13–21
- Vickrey W (1961) Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance* 16:8–37
- Vickrey W (1962) Auctions and bidding games. In: *Recent Advances in Game Theory*, Princeton University Press, Princeton, NJ, no. 29 in Princeton Conference Series, pp 15–27
- Walsh WE, Das R, Tesauro G, Kephart JO (2002) Analyzing complex strategic interactions in multi-agent games. In: *AAAI-02 Workshop on Game Theoretic and Decision Theoretic Agents*
- Walsh WE, Parkes D, Das R (2003) Choosing samples to compute heuristic-strategy nash equilibrium. In: Faratin P, Parkes DC, Rodríguez-Aguilar JA, Walsh WE

-
- (eds) *Agent-Mediated Electronic Commerce V: Designing Mechanisms and Systems*, Springer, Melbourne, Australia, LNAI, vol 3048, pp 109–123, URL cite-seer.ist.psu.edu/walsh03choosing.html
- Watkins JCH, Dayan P (1992) Qlearning. *Machine Learning* 8:279–292
- Weibull JW (1997) *Evolutionary Game Theory*, first mit press edn. MIT Press
- Wellman MP (1995) The economic approach to artificial intelligence. *ACM Computing Surveys* 27(3)
- Wellman MP, Wurman PR, O'Malley K, Bangera R, Lin S, Reeves D, Walsh WE (2001) Designing the market game for a trading agent competition. *IEEE Internet Computing* 5(2):43–51
- Wellman MP, Reeves DM, Lockner KM, Suri R (2005) Searching for walverine. In: *Proceedings of the IJCAI-05 Workshop on Trading Agent Design and Analysis (TADA-05)*, Edinburgh, Scotland
- Williams SR (1988) The nature of equilibria in the buyer's bid double auction. Discussion Paper 793, Department of Economics, Northwestern University
- Williams SR (1991) Existence and convergence of equilibria in the buyer's bid double auction. *The Review of Economic Studies* 58:351–374
- Wolfers J, Zitzewitz E (2004) Prediction markets. *The Journal of Economic Perspectives* 18(2):107–126
- Wurman PR, Walsh WE, Wellman MP (1998) Flexible double auctions for electronic commerce: theory and implementation. *International Journal of Decision Support Systems* 24:17–27
- Wurman PR, Walsh WE, Wellman MP (2001) A parameterisation of the auction design space. *Games and Economic Behaviour* 35:304–338
- Young HP (2001) *Individual Strategy and Social Structure*. Princeton University Press
- Zahavi A, Zahavi A (1997) *The Handicap Principle: A Missing Piece of Darwin's Puzzle*. Oxford University Press
- Zhan W, Friedman D (2005) Markups in double auction markets. Tech. rep., LEEPS, Department of Economics, University of Santa Cruz