



Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 156 (2015) 1–13

JOURNAL OF
**Economic
Theory**

www.elsevier.com/locate/jet

Introduction to computer science and economic theory [☆]

Lawrence Blume ^{a,b,c,*}, David Easley ^a, Jon Kleinberg ^a,
Robert Kleinberg ^a, Éva Tardos ^a

^a *Cornell University, United States*

^b *IHS, Vienna, Austria*

^c *Santa Fe Institute, United States*

Received 3 November 2014; final version received 8 November 2014; accepted 8 November 2014

Available online 18 November 2014

Abstract

This essay introduces the symposium on computer science and economic theory.

© 2014 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-SA license (<http://creativecommons.org/licenses/by-nc-sa/3.0/>).

JEL classification: C7; D6

Keywords: Algorithmic game theory; Implementation; Learning in games; Mechanism design; Networks

1. Introduction

Computer scientists and economists share interests in several areas of economic theory, and individuals from both groups have been working together and in parallel for approximately three decades. Interest in the interaction of computer science and economics has intensified in the last 15 years, fueled in large part by the development of large computer networks such as the Web

[☆] The authors are grateful for support from: L. Blume, NSF grant CCF-0910940, a WWTF grant, an ARO MURI W911NF-12-1-0509 grant; J. Kleinberg, NSF grants IIS-0910664 and CCF-0910940, an ARO MURI W911NF-12-1-0509 grant, a Google Research Grant, and a Simons Investigator Award; R. Kleinberg, NSF grant CCF-0910940, a Google Research Grant, and a Microsoft New Faculty Fellowship Award; É. Tardos, NSF grants CCF-0910940 and CCF-1215994, ONR grant N00014-08-1-0031, a Yahoo! Research Alliance Grant, and a Google Research Grant.

* Corresponding author at: Cornell University, United States.
E-mail address: lb19@cornell.edu (L. Blume).

<http://dx.doi.org/10.1016/j.jet.2014.11.002>

0022-0531/© 2014 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-SA license (<http://creativecommons.org/licenses/by-nc-sa/3.0/>).

and the Internet. The emergence of these new systems has caused a profound expansion of the questions that computer science has had to address, since these networks operate through the cooperation and competition of many participants, leading inevitably to underlying social and economic issues. At the same time, the Internet has also made possible the development of more overtly economic structures, through the creation of new kinds of markets.

The purpose of this symposium is to introduce economists to recent work in these areas. The interaction of computer science and economics has had an impact on economic theory in three ways. It has introduced new problems — novel kinds of markets including those arising in the search industry, and new applications including network management and routing, on-line social systems, and platforms for the production and sharing of content. It has raised new issues in areas already popular in economics, including learning, decision theory, market design, network-structured interaction, the analysis of equilibrium quality, and the computational complexity of equilibria. And it has brought new methods to existing problems, including efficient algorithms, lower bounds based on computational hardness, and techniques from discrete mathematics and graph theory.

We focus in this introduction on the topic of market design as it provides compelling examples of new problems, methods, and techniques in a fundamental economic context, and because mechanism design is at the moment the most active area of joint interest. The next section discusses mechanism design. Subsequent sections provide brief introductions to other aspects of algorithmic game theory, learning in games, and networks.

2. Mechanisms and market design

Marshallian and Walrasian equilibrium analysis are not theories of how markets function. Their institution-free approach to predicting market outcomes precludes them from asking questions such as: When do market institutions fail? How do they behave when they fail? How should markets be designed to minimize failure, and what tradeoffs with market efficiency arise in doing so? Research in economics arising from general equilibrium and welfare economics has been concentrated on market imperfections. Computer scientists have paid relatively more attention to the nuts and bolts of market mechanisms and the robustness of market institutions.

These two distinct approaches to the challenges of modeling markets at a detailed level meet in the field of mechanism design. Leo Hurwicz's research program was a response to the Lange–Lerner–Hayek debate about the virtues of markets versus central planning. To clarify Hayek's claim that the virtue of the market is its ability to harness widely dispersed information to achieve social goals, Hurwicz [1] defined a mechanism to be a communication system in which participants send messages to a center, and a function (or correspondence) which assigns to each profile of messages an allocation of commodities (or a set thereof). Modern mechanism design began with Hurwicz's [2] introduction of *incentive compatibility*, and with Gibbard's [3] introduction of the *revelation principle* for dominant strategy equilibria and its subsequent extension to Bayes–Nash equilibria. Subsequent work has addressed two general questions. The *Implementation problem* is concerned with objectives: Can a particular social choice function be implemented over some rich class of environments? The *Design problem* is concerned with perhaps more practical institutional questions, including: How do particular mechanisms (e.g. first- and second-price auctions) behave? Which auction mechanisms are optimal from a welfare-maximizing point of view, or from a revenue maximizing point of view, or some other point of view? These two questions are not distinct. Although the early implementation theorists (e.g. Hurwicz [4], Mount and Reiter [5]) concluded that markets perform well in “neoclassical” environments,

their abstract description of markets did not enable the analysis of particular methods of market organization, and they had little to say about non-neoclassical environments. These questions are the locus of most mechanism design today. Outside of neoclassical environments, there is yet no general theory of implementation failure, and so the theory develops one problem at a time.

In extending received implementation theory, computer scientists have raised the issue of complexity in mechanism design in two different ways: *Computational complexity*, the difficulty of computing the function that maps message profiles to allocations and of finding an equilibrium, and *communication complexity*, the amount of information that must be communicated by the mechanism's participants. In practical mechanism design both are clear concerns; one needs to have confidence that the algorithm employed in an online market will reach its conclusion in a reasonable amount of time, and market participants need to be able to participate without extensive communication with the mechanism.¹

Computer science has developed a range of ways to measure the efficiency of a procedure or algorithm to solve a problem, based on the resources used by the algorithm; these include the running time, the space or memory required, and the amount of communication required. This also leads to natural measures for the inherent computational complexity of an underlying problem, by considering the minimum resources required by any algorithm to solve the problem. These requirements can be evaluated in either the worst case over all possible initial conditions, or in the average case in a Bayesian environment; for some well-known algorithms, such as the simplex method for linear programming, the difference between worst case and natural measures of average-case performance can be very large.

In a completely analogous way, one can study the amount of computational or communication resources required by a mechanism, and one can define the inherent computational complexity of implementing a social choice function or correspondence as the minimum resources required by any mechanism to implement that function or correspondence. Of particular interest is the question — first raised by Nisan and Ronen [9] — of whether a social choice correspondence which is efficiently computable can always be implemented by a mechanism whose outcome is efficiently computable.² Such a result, if true, would imply that the only obstacles to achieving computational efficiency of mechanisms stem from the inherent computational complexity of the problem being solved, not from the game-theoretic challenge of implementing the solution in equilibrium. Unfortunately, the result is false, as can be seen by considering the problem of implementing social welfare maximization with quasi-linear utility in dominant-strategy equilibrium. The Vickrey–Clarke–Groves (VCG) mechanism is unique for implementing social welfare maximization, and requires solving an NP-hard problem (a widely adopted criterion for computational intractability) in settings with an exponentially large discrete set of alternatives, as was observed by Kfir-Dahav, Monderer, and Tennenholtz [10].

In light of this negative result, attention naturally turned to implementing outcomes that achieve close to maximal welfare. The answer to whether the computational complexity of *implementing* this social choice correspondence differs significantly from the complexity of merely *computing* it depends on the details of the problem. In some cases there exist computationally efficient algorithms implementing an outcome whose efficiency is within a small constant factor of the maximal social welfare. In other cases, computational constraints can be arbitrarily

¹ The issue of computational complexity has been raised in economic contexts, motivated by CS research; for example, by [6–8].

² Here, as in many contexts in computer science, efficient computability is defined by stipulating that the quantity of resources required for a computation must be bounded above by a polynomial function of its input size.

costly for the efficiency of outcome. An example of this latter case is the work of Papadimitriou, Schapira and Singer [11] on “combinatorial public projects”, that is, selecting k out of n potential public projects to maximize the combined utility of n bidders with submodular valuations over sets of projects. Selecting an outcome that achieves fraction $1 - 1/e$ of the maximal social welfare can be achieved quite easily by polynomial-time algorithms, but they show that mechanisms implementing this social choice correspondence in dominant-strategy equilibrium cannot be computationally efficient. The same phenomenon, whereby incentive-compatibility constraints lead to an exponential blow-up in the computational resources required to solve a problem, occurs even for the seemingly innocuous problem of allocating identical, indivisible goods to a set of bidders, as shown by Dobzinski and Nisan [12].

This strand of research illustrates a methodological innovation of computer scientists. If utility is ordinal or cardinal but not interpersonally comparable, then implementation results are of the form, “mechanism x does (does not) achieve a pareto optimum”. If utility is cardinal and interpersonally comparable, then aggregate welfare is cardinal, and it becomes meaningful to say things like, “mechanism x achieves at least $1 - 1/e$ of maximal social welfare”. This allows for a much finer gradation of mechanisms that are not socially optimal — and since most problems do not admit optimal mechanisms with a tractable structure, the ability to discuss aggregate welfare opens up many new and interesting questions.

Communication complexity asks the question, “How much information needs to be exchanged between a set of parties who are collectively carrying out a computation?” This question has its origins in computer science [13,14]. The following example comes from Nisan and Segal [15], a collaboration between a computer scientist and an economist. Combinatorial auctions are mechanisms designed to allocate L heterogeneous items among N bidders whose valuations for the different objects are not necessarily additive across the items. The revelation principle reduces this problem to that of studying revelation games, but full revelation of a bidder’s preferences requires transmission of a willingness to pay for each of $2^L - 1$ bundles of items. That is over 1 billion numbers for $L = 30$. Nisan and Segal provide results both for bits and, for those cases involving continuous information, dimension. Dobzinski and Nisan provide a useful illustration of the types of methods that enable analyzing the communication requirements of mechanism design problems.

This issue of communication complexity was at the heart of the Hayek–Lange–Lerner debate, and for that reason, among many others, it should be familiar to the economics community. Hayek [16, p. 519] argued that: “... the ‘data’ from which the economic calculus starts are never for the whole society ‘given’ to a single mind which could work out the implications and can never be so given”. The “economic problem of society” is “a problem of the utilization of knowledge which is not given to anyone in its totality”. Subsequently, Hurwicz [4] and Mount and Reiter [5] demonstrated that Hayek was correct for neoclassical environments, that is, resource allocation problems with no externalities and diffuse market power. Markets are “informationally efficient”, that is, they minimize the amount of communication needed to support an optimal allocation of resources, where the quantity of communication is measured by the dimensionality of the space of signals needed to convey all necessary information. Saari and Simon [17] showed that the amount of information needed to find Walrasian prices is nonetheless large. Global Newton methods will (generically) find zeros of excess demand functions. For an $n + 1$ -commodity economy they require at least $n^2 + n$ pieces of information: the n independent values of excess demand and their n^2 derivatives with respect to normalized prices. No method (in the sense that the global Newton method is a method) that works on any excess demand function can do better when $n \geq 3$.

Early on in the mechanisms literature it was recognized that dominant strategy implementability was too strong — little could be implemented. The economics mechanism-design community quickly took up Bayesian implementation, following [18], Harris and Townsend [19], Holmström [20], and Myerson [21]. Computer science has considered Bayesian implementations that achieve approximately some design criteria (like maximizing welfare or revenue). But they also have been more aggressive in looking for other, non-Bayesian alternatives to dominant strategies. Lavi and Nisan [60] is an example of this work. They investigate a dynamic mechanism design problem using what they call “set-Nash equilibrium”: Each player is given a recommended set of strategies and may play an arbitrary element of the recommended set, but the n -tuple of recommended strategy sets has the stability property that if every other player chooses an element of their recommended set, then the left-over player must have a best response in her recommended set. Chen and Micali [61] propose an alternate approach to weakening underlying assumptions in auctions, using a purely set-theoretic model of beliefs.

More broadly, a perspective that guides a lot of work in computer science is to achieve approximately desired outcomes under minimal assumptions about the state of the world. This perspective leads naturally to questions about alternatives to Bayesian analysis, which makes strong assumptions about distributions of problem inputs (e.g. valuations in an auction) and looks to maximize average social welfare or average seller revenue. Robustness aside, Bayesian games are limited in application to the degree they rely on common knowledge assumptions. Robert Wilson [22] has famously written, “I foresee the progress of game theory as depending on successive reductions in the base of common knowledge required to conduct useful analysis of practical problems. Only by repeated weakening of common knowledge assumptions will the theory approximate reality”. Computer science perspectives can thus be viewed by economists as working in alignment with what has come to be called Wilson’s Doctrine.

One approach to mechanism design with reduced assumptions is a “prior-independent mechanism”, that is, a detail-free mechanism whose designer does not need to know things like type distributions. Such results are situated halfway between average-case and worst-case analysis: For any given distribution of types, the mechanism’s performance is evaluated in the average case (that is, expected performance with respect to the *a priori* type distribution) but the theorem that provides the performance guarantee for the mechanism holds in the worst case over all distributions. One sees this type of goal in early examples such as Bulow and Klemperer’s [23] result that a second-price auction with $n + 1$ i.i.d. bidders will get at least as much revenue, in expectation, as an optimal auction with n bidders from the same distribution. (Note that running a second-price auction doesn’t require any knowledge of bid distributions, so this is an example of a detail-free result.) Results in this style have been given broad generalizations by work in computer science, including results of Hartline and Roughgarden [24], and new results of Dhangwatnotai, Roughgarden and Yan [25] who consider a class of environments including single-unit auctions, k -unit auctions and more, with private and independent bidder valuations drawn from monotone hazard rate distributions which are common knowledge among the bidders, but not necessarily the same. They construct a mechanism that, regardless of the prior, guarantees an expected revenue to the seller of at least $1/2$ the expected revenue achievable by the optimal mechanism for that environment.

An even more aggressive weakening of Bayesian assumptions takes place in the “prior-free” approach to mechanism design, which aims to provide meaningful worst-case revenue guarantees for mechanisms, i.e. non-trivial lower bounds on a mechanism’s revenue that hold for every profile of bids and not just in expectation. This approach is exemplified by the paper of Devanur et al. [62] that develops a framework for prior-free mechanism design and analysis in

single-dimensional settings. They provide meaningful and non-trivial revenue lower bounds for a natural random-sampling-based mechanism by comparing its revenue on every bid profile to a benchmark value defined by optimizing over outcomes that are “envy-free” with respect to the same bid profile.

Computer science research has also studied these weakenings of assumptions to ask questions about “what is” rather than about “what is optimal”. Myerson’s optimal mechanism is beautiful, and it is indeed simple when agents have identical type distributions. For non-identical distributions, it differs from the simple auction formats seen in reality (which use anonymous reserve prices, for example). Hartline and Roughgarden [24] and subsequent work have proposed an appealing quantitative argument for the prevalence of simple auctions, showing that they achieve a bounded approximation factor with respect to the revenue of optimal auctions when bidders’ types are single-dimensional and drawn from distributions satisfying a natural regularity condition. Briest et al. [63] provide a contrasting result in a domain with multi-dimensional preferences. They focus on one aspect of simple auctions — that their outcome is a deterministic function of the bid profile — and show how the use of randomization in certain settings can achieve unbounded gains in revenue relative to the best deterministic mechanisms. This research program, and subsequent work of Haghpanah and Hartline [26] are inverse results to the classical implementation problem, whereas implementation theorists (e.g. Maskin [27]) ask for which social choice functions do there exist mechanisms which can implement them,³ this program asks for which environments a given mechanism implements or approximately implements a given social choice function, in this case revenue maximization.

Although the analysis of market institutions is associated with the mechanism design of the 1960s and 1970s, there has been independent interest in institutional analysis almost since the invention of game theory.⁴ Neoclassical analysis analyzes market outcomes by looking at the consequences of small (marginal deviations) and making assumptions (such as convexity) that allow for the inference of global properties from local properties. Theorists such as Gale, Shapley, Shubik and Scarf, on the other hand, foresaw that indivisibilities are not just a technical problem. Distinct from the marginal tradition of neoclassical economics, they saw the problem of optimal resource allocation as a matching problem. Gale and Shapley [29], Shapley and Scarf [30], and Shapley and Shubik [31] introduced one- and two-sided matching. This distinction becomes particularly important in two-sided matching markets, such as matching students to schools and workers to firms, and represents an area of mechanism design that has had a significant influence on empirical economics. Matching markets and related design problems such as kidney exchanges are an active area of current research; a formulation of matching problems involving multiparty contracts is considered by the paper of Hatfield and Kominers [64], who show welfare and stability results for such multilateral interactions.

3. Computation in games

Moving beyond mechanism design, computational complexity is important more generally for game theory. Recent computational work [32,33] has identified fundamental limits on the ability of efficient algorithms to compute equilibria. These results show how to take very hard instances of computational problems and embed them in a set of strategies and payoffs for players in a

³ Strongly, in the sense that *all* equilibria achieve social choice-function outcomes.

⁴ In fact, even earlier: the first formal analysis of a game was Beccaria’s [28] optimal tariff problem.

game — with the consequence that if the players are able to reach an equilibrium, they would collectively produce a solution to problems that we believe to be computationally infeasible.

This suggests some inherent limits on the power of equilibrium notions — that for certain difficult games, we should not expect players to be able to find an equilibrium. However, the kinds of games used to establish these results have a fragile, pathological flavor to them — based on all-or-nothing situations in which one must essentially find the “secret combination” to solve a puzzle. In contrast, games at the heart of market behavior and other forms of real-world strategic interaction seem to be smoother and more robust — despite the abundance of sub-optimal outcomes, the system tends to contain cues that guide the behavior of the participants, and forms of feedback that can select for better behavior. A broad research question is to quantify the features of games in real settings that lead to more tractable behavior. In a related direction, one can ask which natural classes of games do not suffer from the difficulty of computing equilibria. For example, Daskalakis and Papadimitriou [65] identify the class of anonymous games where they show that a minimal relaxation of equilibrium conditions eases the difficulty of computing equilibria.

Complexity also has implications for the modeling of boundedly rational agents. One source of bounded rationality is limited computational resources. Neyman [34] and Rubinstein [35] modeled bounded rationality in finitely repeated prisoners’ dilemmas by requiring agents to employ strategies that could be implemented by finite state automata. If the size of the automata (the number of states) is sufficiently constrained, Nash equilibria in automata strategies can support cooperation. With enough states, though, Papadimitriou and Yannakakis [36] showed that finite automata can compute best responses to every partial history, and so the only Nash equilibrium in such automata strategies is to always defect. Halpern and Pass [66] provide new results on the use of costly computation as a model of bounded rationality. Rather than limiting players to computations that can be implemented using finite automata, they model fully general computation using Turing machines, and this motivates alternative ways of quantifying the cost of a computation, such as the number of random bits consumed.

4. The price of anarchy and quality of learning outcomes

The Lange–Lerner–Hayek debate had to do with the relative advantages and disadvantages of decentralized mechanisms. Although the market proves to be efficient for neoclassical environments, this is not necessarily the case in other environments. But the issue is broader than just the consequences of decentralized information — for games in general one can ask how well an equilibrium outcome does relative to the social optimum. Koutsoupas and Papadimitriou (1999) first formalized this idea in the context of Nash equilibria by computing the *price of anarchy*, the ratio of the welfare of the socially optimal outcome of a game to the minimum welfare of any Nash equilibrium outcome. Note that this line of research again illustrates the types of questions that can be asked when one views utilities as cardinal and interpersonally comparable, allowing aggregate measures of welfare to be evaluated.

The price of anarchy has proven to be a useful framework for analyzing a number of settings where self-interested behavior leads to sub-optimal outcomes. An early influential result comes from the work of Roughgarden and Tardos [37] on network congestion. Suppose that travel time on a road increases with the degree of congestion. How would one allocate drivers to different routes on the network so as to minimize average travel time? A centralized solution would assign a particular route to each driver. A decentralized solution has each individual choosing the time-minimizing route given the route choices of others. Among other results, Roughgarden and

Tardos show that the total welfare of a decentralized solution with strategic participants is no worse than a centralized optimal solution with twice the level of traffic.

An important limitation of the price of anarchy is that it assumes that the players have successfully coordinated on a Nash equilibrium, which can be a strong assumption. However, many results on the price of anarchy turn out to be rather robust in that they extend to a form of equilibrium that players reach via learning. Many situations give rise to repeated series of strategic interactions, that is, repeated games. Routing is such a game, as are many “small stakes” auctions, such as sponsored search auctions. In these situations it is interesting to study not only the set of equilibria, but how they are achieved when players refine their play as they learn more about other participants and even the rules of the game. In contrast to “high-stakes games” where it makes sense to invest significant resources in learning to play well up front, such investments are not worthwhile if the stakes in any single interaction are sufficiently low. Nonetheless we expect players to learn, and the long-run consequences of learning are significant when the volume of interactions is high. Learning in games has a long history, going back to early work on fictitious play [38].⁵ Much of this work models Bayesian learning, in which case the relevant solution concept, Markov perfect equilibrium, is notoriously complex. Adlakha et al. [67] explore the use of an alternative *stationary equilibrium* solution concept; they provide conditions guaranteeing that stationary equilibria exist and closely approximate Markov perfect equilibria, and they explore the consequences for the phenomenon of “learning by doing” in an oligopoly model.

An alternative to Bayesian learning that plays an important role in the computer science literature is “no-regret learning”. In a repeated game we suppose that in each stage a player observes the consequences of all his possible actions given the actual play of his opponents. The *regret* of the action sequence a_1, \dots, a_T of length T is the excess of the payoff from the best *fixed* action in hindsight to the payoff realized by the sequence. A learning algorithm is just a strategy for the repeated game — a function mapping histories of play into actions. A learning algorithm is *no-regret* if regret is guaranteed to grow sublinearly in T , no matter how opponents play. No-regret algorithms exist, and they must necessarily randomize. Furthermore, no-regret learning converges to an equilibrium in zero-sum games, and more generally to a weakening of correlated equilibrium called *coarse correlated equilibrium*, a correlated distribution of actions such that no player would prefer to switch to any fixed action rather than play the given correlated distribution. All correlated equilibria (and hence all Nash equilibria) are coarse correlated equilibria.

An interesting recent development using the idea of learning as a model of game outcomes is that in many games the known price of anarchy results extends to learning outcomes, i.e., the welfare guarantee provided by the price of anarchy result applies even when no-regret-learning players have not converged to a Nash equilibrium. A compelling result on this theme is due to Roughgarden [40], who demonstrates that if a game satisfies a particular payoff condition which he calls a “smoothness condition”, then price-of-anarchy results for that game extend to coarse correlated equilibria, the limits of no-regret learning. The proposed smoothness condition is satisfied by most games with known price of anarchy bounds, including the routing game and games arising in sponsored search auctions. The paper of Roughgarden and Schoppmann [68] weakens the smoothness condition to extend it to a class of games with convex strategy sets, while the weaker condition still implies that the price-of-anarchy results for that game extend to correlated equilibria (though not to coarse correlated equilibria).

⁵ A recent survey is Hart and Mas-Colell [39].

5. Applied market design

Economists and computer scientists have also collaborated productively in actually designing markets and in analyzing these designed markets. A leading example is the sponsored search market. Web search companies such as Google and Yahoo! raise billions of dollars annually by selling space beside particular search queries. Current on-line advertising markets look very different from the advertising markets of earlier eras; whereas sales in these earlier forms of the market were determined by individual negotiations, the on-line versions involve an extremely rich structure for the items being sold, based on the attention of individual users as they seek to accomplish specific search tasks. In analyzing and designing such markets, there are opportunities to exploit connections to some of the fundamental models of matching markets that have been the subject of extensive study in both computer science and economics.

The first papers to propose and analyze a model of sponsored search were Edelman et al. [41], Varian [42], and Mehta et al. [43]. Edelman et al. and Varian show that an efficient equilibrium for basic models of sponsored search always exists in the full information setting. The paper of Caragiannis et al. [69] considers this framework for a Bayesian model of sponsored search auctions (specifically, generalized second-price auctions with uncertain quality scores), and shows a small constant bound on the price of anarchy that is robust enough to extend to the incomplete information game even with correlated types (and correlated quality scores).

Prediction markets form another broad class of markets that have become much more prevalent with powerful computational resources and the growth of the Internet [44]. A prediction market is one that is designed for the purpose of inducing a collective prediction from the set of participants — for example, awarding a dollar if a particular candidate is elected, or a particular team wins a sporting event (e.g. [45–47]). The notion draws on an idea with a long history in economics, that prices reveal beliefs; with greater computational power and lightweight access to large groups of potential participants through the Internet, prediction markets have become correspondingly more expressive, with award conditions that now extend to complex Boolean predicates on the basic outcomes. A central goal of a prediction market is to be able to sell and buy securities guaranteeing good prediction while limiting the cost to the market maker, and developing computationally efficient pricing mechanisms has been an active focus of research in the computer science community. Running prediction markets online gives rise to some other novel issues as well. For example, good online mechanisms need to be sybilproof, meaning that individuals should not be able to benefit from submitting several reports under different pseudonyms, as such sybils are easy to generate online. Lambert et al. [70] gives an axiomatic characterization of mechanisms for prediction markets satisfying a number of desirable criteria.

6. Networks

Networks have been a basic object in computer science since its initial stages, as the formalism of graph theory has proved to be a powerful language for modeling many of the field's core applications beginning with some of the earliest: wiring diagrams of circuits, control flow in programs, the structure of communication and transportation networks, and then more recently the topology of the World Wide Web and large social media and on-line commerce sites. These more recent networks have increasingly consisted of interacting strategic agents, rather than components that were centrally designed by a single entity, and so it became natural for economic and

game-theoretic considerations to merge increasingly with the computational considerations. Economics also began to investigate graphs as the underlying models for some of its core questions; these include the interactions of participants in a market, and questions concerning the diffusion of information, innovations, and behavior in an underlying social network.

Both fields have considered how a network's topology can arise as the result of strategic interactions among the nodes; Shenker [48] asked how such considerations would affect the structure of the growing Internet topology, and work beginning with Jackson and Wolinsky [49] in economic theory and Fabrikant et al. [50] and [51] in the theory of computing studied the consequences of self-interested behavior in the formation of network structure.

In addition to the structure of the network itself, there has been significant interest in both computer science and economics on the processes that take place on the underlying network. The routing of traffic on networks is one such process, which as noted earlier formed one of the core settings for analyzing the price of anarchy [37]. Other active lines of work include analyses of the diffusion of information, innovations, and behaviors on social networks (Rogers [52], Blume [53], and Ellison [54]); and the outcome of buyer–seller interaction and bargaining on network structures (Kranton and Minehart [55], Corominas-Bosch [56], and Kakade et al. [57]).

The perspective of computer science is evident in recent approaches to both of these two latter issues. In the diffusion of innovations, one is led to an interesting algorithmic variation on the underlying question by considering how to intervene in it — in particular, how a centralized agent could optimally “seed” an innovation through careful selection of the starting nodes. This was asked by Domingos and Richardson [58]; since the optimal choice of seeds is computationally difficult, and lacking in any apparent tractable structure, subsequent work identified rich structural properties of near-optimal solutions that can be found through greedy optimization, by establishing a submodular property of the process (Kempe, Kleinberg, and Tardos [59]). In the context of trading and bargaining on networks, where participants in the network may repeatedly update their offers to others, there are deep connections to distributed updating algorithms related to *belief propagation*; Bayati et al. [71] develop this connection to provide new results on networked bargaining.

7. Conclusion

There is something natural in the broadening interface between computer science and economics — each can be viewed as a field focused on the design and analysis of extremely complex synthetic systems that are governed by phenomena we only partially understand and only partially know how to control. And each field increasingly needs knowledge that resides in the other, as computer scientists seek to design and understand systems containing ever-increasing numbers of strategic agents, and economists seek to model the complex networks of interactions that lie between small-group environments and population-scale economic activity.

This introduction has been designed to serve as a roadmap to some of the central themes in the papers that follow. The questions that these papers address in turn point to important issues that will continue to shape the interface between computer science and economics: establishing guarantees for economic processes under increasingly general assumptions; developing methods to evaluate the performance of algorithms and mechanisms even in cases when they may not be the optimal choice; delineating the ways in which computational efficiency serves both to make certain outcomes feasible and to provide evidence for the infeasibility of others; and identifying design principles for the complex environments that arise when strategic agents come together, make decisions, and interact.

References

- [1] L. Hurwicz, Optimality and informational efficiency in resource allocation processes, in: K.J. Arrow, S. Karlin, P. Suppes (Eds.), *Mathematical Methods in the Social Sciences*, Stanford University Press, Palo Alto, CA, 1960.
- [2] L. Hurwicz, On informationally decentralized systems, in: C.B. Maguire, R. Radner (Eds.), *Decision and Organization*, North-Holland, Amsterdam, 1972.
- [3] A. Gibbard, Manipulation of voting schemes: a general result, *Econometrica* (1973) 587–601.
- [4] L. Hurwicz, On the dimensional requirements of informationally decentralized pareto-satisfactory processes, in: K.J. Arrow, L. Hurwicz (Eds.), *Studies in Resource Allocation Processes*, Cambridge University Press, New York, 1977, pp. 413–424.
- [5] K. Mount, S. Reiter, The informational size of message spaces, *J. Econ. Theory* 8 (1974) 161–192.
- [6] S.E. Spear, Learning rational expectations under computability constraints, *Econometrica* (1989) 889–910.
- [7] L. Anderlini, H. Sabourian, Cooperation and effective computability, *Econometrica* 63 (1995) 1337–1369.
- [8] L. Anderlini, Communication, computability, and common interest games, *Games Econ. Behav.* 27 (1999) 1–37.
- [9] N. Nisan, A. Ronen, Algorithmic mechanism design, *Games Econ. Behav.* 35 (2001) 166–196; preliminary version appeared in: *ACM Symposium on the Theory of Computing*, ACM, 1999, pp. 129–140.
- [10] N.E. Kfir-Dahav, D. Monderer, M. Tennenholtz, Mechanism design for resource bounded agents, in: *Proceedings. Fourth International Conference on MultiAgent Systems*, IEEE, 2000, pp. 309–315.
- [11] C. Papadimitriou, M. Schapira, Y. Singer, On the hardness of being truthful, in: *Forty-Ninth Annual IEEE Symposium on Foundations of Computer Science*, 2008, pp. 250–259.
- [12] S. Dobzinski, N. Nisan, Multi-unit auctions: beyond Roberts, *J. Econ. Theory* 156 (2015) 14–44.
- [13] A.C.-C. Yao, Some complexity questions related to distributive computing (preliminary report), in: *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing*, ACM, 1979, pp. 209–213.
- [14] E. Kushilevitz, Communication complexity, *Adv. Comput.* 44 (1997) 331–360.
- [15] N. Nisan, I. Segal, The communication requirements of efficient allocations and supporting prices, *J. Econ. Theory* 129 (2006) 192–224.
- [16] F.A. Hayek, The use of knowledge in society, *Amer. Econ. Rev.* (1945) 519–530.
- [17] D.G. Saari, C.P. Simon, Effective price mechanisms, *Econometrica* (1978) 1097–1125.
- [18] P. Dasgupta, P. Hammond, E. Maskin, The implementation of social choice rules: some general results on incentive compatibility, *Rev. Econ. Stud.* (1979) 185–216.
- [19] M. Harris, R.M. Townsend, Resource allocation under asymmetric information, *Econometrica* (1981) 33–64.
- [20] B.R. Holmström, On incentives and control in organizations, PhD thesis, Stanford University, Palo Alto, CA, 1977.
- [21] R.B. Myerson, Incentive compatibility and the bargaining problem, *Econometrica* (1979) 61–73.
- [22] R. Wilson, Game-theoretic analyses of trading processes, in: T.F. Bewley (Ed.), *Advances in Economic Theory*, Cambridge University Press, 1987, pp. 33–70.
- [23] J. Bulow, P. Klemperer, Auctions versus negotiations, *Amer. Econ. Rev.* (1996) 180–194.
- [24] J.D. Hartline, T. Roughgarden, Simple versus optimal mechanisms, in: *Proceedings of the Tenth ACM Conference on Electronic Commerce*, 2009, pp. 225–234.
- [25] P. Dhangwatnotai, T. Roughgarden, Q. Yan, Revenue maximization with a single sample, in: *Proceedings of the Eleventh ACM Conference on Electronic Commerce, EC '10*, ACM, 2010, pp. 129–138.
- [26] N. Haghpanah, J. Hartline, Reverse mechanism design, arXiv preprint, arXiv:1404.1341, 2014.
- [27] E. Maskin, Nash equilibrium and welfare optimality, MIT, 1977, unpublished.
- [28] C. Beccaria, Tentativo analitico su i contrabbandi, in: G. Francioni, S. Romagnoli (Eds.), *Il Caffè*, Bollati, Torino, 1766, 1993.
- [29] D. Gale, L.S. Shapley, College admissions and the stability of marriage, *Amer. Math. Monthly* 69 (1962) 9–14.
- [30] L. Shapley, H. Scarf, On cores and indivisibility, *J. Math. Econ.* 1 (1974) 23–37.
- [31] L.S. Shapley, M. Shubik, The assignment game I: the core, *Int. J. Game Theory* 1 (1971) 111–130.
- [32] C. Daskalakis, P.W. Goldberg, C.H. Papadimitriou, The complexity of computing a Nash equilibrium, in: *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, 2006, pp. 71–78.
- [33] X. Chen, X. Deng, Settling the complexity of two-player Nash equilibrium, in: *Forty-Seventh Annual IEEE Symposium on Foundations of Computer Science*, vol. 6, 2006.
- [34] A. Neyman, Bounded complexity justifies cooperation in the finitely repeated prisoners' dilemma, *Econ. Letters* 19 (1985) 227–229.
- [35] A. Rubinstein, Finite automata play the repeated prisoner's dilemma, *J. Econ. Theory* 39 (1986) 83–96.

- [36] C.H. Papadimitriou, M. Yannakakis, On complexity as bounded rationality, in: Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, ACM, 1994, pp. 726–733.
- [37] T. Roughgarden, É. Tardos, How bad is selfish routing?, *J. ACM* 49 (2002) 236–259; preliminary version appeared in: Annual IEEE Symposium on the Foundations of Computer Science, IEEE, 2000, pp. 93–102.
- [38] G.W. Brown, Iterative solution of games by fictitious play, in: T.C. Koopmans (Ed.), *Activity Analysis of Production and Allocation*, Wiley, New York, 1951, pp. 374–376.
- [39] S. Hart, A. Mas-Colell, Y. Babichenko, *Simple Adaptive Strategies: From Regret-Matching to Uncoupled Dynamics*, vol. 4, World Scientific Publishing Company, 2013.
- [40] T. Roughgarden, Intrinsic robustness of the price of anarchy, in: Proceedings of the 41st Annual ACM Symposium on Theory of Computing, 2009, pp. 513–522.
- [41] B. Edelman, M. Ostrovsky, M. Schwarz, Internet advertising and the generalized second-price auction: selling billions of dollars worth of keywords, *Amer. Econ. Rev.* 97 (2007) 242–259.
- [42] H.R. Varian, Position auctions, *Int. J. Ind. Organ.* 26 (2007) 1163–1178.
- [43] A. Mehta, A. Saberi, U. Vazirani, V. Vazirani, Adwords and generalized online matching, *J. ACM* 54 (2007) 22; preliminary version appeared in: Annual IEEE Symposium on Foundations of Computer Science, 2005.
- [44] D.M. Pennock, R. Sami, Computational aspects of prediction markets, in: N. Nisan, T. Roughgarden, É. Tardos, V. Vazirani (Eds.), *Algorithmic Game Theory*, Cambridge University Press, New York, 2007, pp. 651–674.
- [45] M.M. Ali, Probability and utility estimates for racetrack bettors, *J. Polit. Economy* 85 (1977) 803–815.
- [46] R.H. Thaler, W.T. Ziemba, Anomalies: parimutuel betting markets: racetracks and lotteries, *J. Econ. Perspect.* 2 (1988) 161–174.
- [47] R. Forsythe, F. Nelson, G.R. Neumann, J. Wright, Anatomy of an experimental political stock market, *Amer. Econ. Rev.* (1992) 1142–1161.
- [48] S. Shenker, Fundamental design issues for the future Internet, *IEEE J. Sel. Areas Commun.* 13 (1995) 1176–1188.
- [49] M.O. Jackson, A. Wolinsky, A strategic model of social and economic networks, *J. Econ. Theory* 71 (1996) 44–74.
- [50] A. Fabrikant, A. Luthra, E. Maneva, C.H. Papadimitriou, S. Shenker, On a network creation game, in: Proceedings of the Twenty-Second Annual Symposium on Principles of Distributed Computing, PODC '03, 2003, pp. 347–351.
- [51] E. Anshelevich, A. Dasgupta, E. Tardos, T. Wexler, Near-optimal network design with selfish agents, in: Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing, STOC '03, 2003, pp. 511–520.
- [52] E.M. Rogers, *Diffusion of Innovations*, Free Press, New York, 1962.
- [53] L.E. Blume, The statistical mechanics of strategic interaction, *Games Econ. Behav.* 5 (1993) 387–424.
- [54] G. Ellison, Learning, local interaction and coordination, *Econometrica* 61 (1993) 1047–1071.
- [55] R.E. Kranton, D.F. Minehart, A theory of buyer–seller networks, in: *Networks and Groups*, Springer, 2003, pp. 347–378.
- [56] M. Corominas-Bosch, Bargaining in a network of buyers and sellers, *J. Econ. Theory* 115 (2004) 35–77.
- [57] S.M. Kakade, M. Kearns, L.E. Ortiz, R. Pemantle, S. Suri, Economic properties of social networks, in: *Advances in Neural Information Processing Systems*, 2005, pp. 633–640.
- [58] P. Domingos, M. Richardson, Mining the network value of customers, in: Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '01, 2001, pp. 57–66.
- [59] D. Kempe, J. Kleinberg, É. Tardos, Maximizing the spread of influence through a social network, in: Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2003, pp. 137–146.
- [60] R. Lavi, N. Nisan, Online ascending auctions for gradually expiring items, *J. Econ. Theory* 156 (2015) 45–76.
- [61] J. Chen, S. Micali, Mechanism design with possibilistic beliefs, *J. Econ. Theory* 156 (2015) 77–102.
- [62] N.R. Devanur, J.D. Hartline, Q. Yan, Envy freedom and prior-free mechanism design, *J. Econ. Theory* 156 (2015) 103–143.
- [63] P. Briest, S. Chawla, R. Kleinberg, S.M. Weinberg, Pricing lotteries, *J. Econ. Theory* 156 (2015) 144–174.
- [64] J.W. Hatfield, S.D. Kominers, Multilateral matching, *J. Econ. Theory* 156 (2015) 175–206.
- [65] C. Daskalakis, C.H. Papadimitriou, Approximate Nash equilibria in anonymous games, *J. Econ. Theory* 156 (2015) 207–245.
- [66] J.Y. Halpern, R. Pass, Algorithmic rationality: game theory with costly computation, *J. Econ. Theory* 156 (2015) 246–268.
- [67] S. Adlakha, R. Johari, G.Y. Weintraub, Equilibria of dynamic games with many players: existence, approximation, and market structure, *J. Econ. Theory* 156 (2015) 269–316.
- [68] T. Roughgarden, F. Schoppman, Local smoothness and the price of anarchy in splittable congestion games, *J. Econ. Theory* 156 (2015) 317–342.
- [69] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, B. Lucier, R. Paes Leme, É. Tardos, Bounding the inefficiency of outcomes in generalized second price auctions, *J. Econ. Theory* 156 (2015) 343–388.

- [70] N.S. Lambert, J. Langford, J.W. Vaughan, Y. Chen, D.M. Reeves, Y. Shoham, D.M. Pennock, An axiomatic characterization of wagering mechanisms, *J. Econ. Theory* 156 (2015) 389–416.
- [71] M. Bayati, C. Borgs, J. Chayes, Y. Kanoria, A. Montanari, Bargaining dynamics in exchange networks, *J. Econ. Theory* 156 (2015) 417–454.