Electronic Commerce: From Economic and Game-Theoretic Models to Working Protocols

Moshe Tennenholtz
Faculty of Industrial Engineering and Management
Technion - Israel Institute of Technology
Haifa 32000, Israel

1 Introduction

The area of mechanism design is fundamental to economics and game theory. Research in this area is concerned with the design of multi-agent interactions in ways that lead rational agents to desired behavior [Fudenberg and Tirole, 1991; Kreps, 19901. The design of protocols is fundamental to Computer Science, and is concerned with the design of multi-agent interactions in computational settings, in order to satisfy a given system specification. At first glance, it seems that mechanism design and protocol design deal with the same issues, but a further study shows that these areas have been concerned with different aspects of multi-agent interactions. Work in mechanism design has been mostly concerned with rational behavior in non-cooperative systems, where the agents' utilities play a major role. Work on the design of protocols has been concerned with multi-agent interactions in distributed systems, where issues such as security, network topology, fault tolerance, and formal specification and verification play a major role. In addition, the mathematical techniques used in these areas have been vastly different. While real analysis is the standard tool for the study of economic mechanisms, logic, combinatorics, and graph theory are the standard tools for the analysis of protocols. This situation leaves us with a significant challenge: in order to design protocols to be used in non-cooperative computational environments, such as the Internet, all aspects of mechanism and protocol design should be integrated into a coherent theory.

The transition from economic and game-theoretic models to working protocols is a major challenge for electronic commerce. In order to design useful economic mechanisms for computational settings, a joint effort of experts in economics, game theory, computer science, and artificial intelligence is essential. We believe that the Internet will motivate new theories in Economics and Computer Science, and that the design of protocols/mechanisms based on the these new theories will lay foundations to a theory of electronic commerce.

This report introduces a (biased) perspective on the relationships between economic mechanisms and computational settings/protocols. In Section 2 we briefly discuss the theory of mechanism design, emphasizing two basic settings which illustrate the basic structure of that theory. In Section 3 we show how insights obtained by the implementation of economic mechanisms in computational environments, lead to new basic economic models, questions, and solutions. In Section 4-6 we discuss several issues in the transition from economic mechanisms to computational protocols; in particular, we discuss the role played by the communication graph, by the actual representation of the agents' information and messages, by the asynchronous nature of distributed systems, and by the agents' resource bounds, when transforming mechanisms into working protocols.

2 Mechanism Design

Work in mechanism design deals with the design of multi-agent interactions of rational agents. This work is fundamental to economics [Kreps, 19901, as well as to the attempt to design protocols for non-cooperative computational environments [Boutilier *ex al.,* 1997; Rosenschein and Zlotkin, 1994; Bond and Gasser, 1988]. The importance of the latter has increased dramatically due to the desire to design economic interactions in the Internet setup. In such a context, the fact we have different human users in a distributed system, suggests we can not rely on standard protocol design; we have to design protocols that are incentive-compatible. In order to handle this problem, the theory of mechanism design has to be adapted to computational environments, i.e. to the context of distributed computing.

In a typical mechanism design setup [Fudenberg and Tirole, 1991; Kreps, 1990; Monderer and Tennenholtz, 1998b], a center attempts to perform some task, or to arrive at some decision, based on information available to a set of agents. The major problem is that the agents might not supply the desired information, and might cheat in order to increase their individual payoffs. We therefore need to design a mechanism that when the agents interact rationally through that mechanism, the center will be able to obtain its objective. For example, the center may wish to sell a particular good while obtaining maximal profit, or to distribute goods among the agents in a way that maximizes social welfare. In order to do so, a mechanism has to be designed; typically, this mechanism determines actions to be taken by the center as a function of messages sent to it by the agents; the agents' strategies correspond to the messages they send to the center as a function of the private information they have. Such interaction is modeled by a Bayesian game [Fudenberg and Tirole, 1991].

The aim is to design a mechanism, such that when the agents behave rationally, i.e. according to an equilibrium of the corresponding game, a desired behavior is obtained.

It will be impossible to present a detailed overview of the theory of mechanism design in the framework of this short paper. Instead, we will describe two fundamental settings which illustrate the basic ideas of mechanism design.

2.1 Mechanisms for Optimal Revenue

Our first representative mechanism design setting refers to a situation where a designer wishes to maximize its revenue from an economic transaction. Auctions [Wolfstetter, 1996; Miigrom, 1987; Wilson, 1992] are the most typical economic mechanisms in this regard,

Assume a center/seller wishes to sell a good/service to one of *n* agents. The private information of each agent refers to its valuation of the good, Le. its maximal willingness to pay for the good. The assumption is that the agents' valuations for the good are independently drawn from a given interval (e.g. the valuation for a particular TV is in the range between $100 and $500), according to a given distribution function (e.g. the uniform distribution). Each agent knows its valuation for the good, but does not know the valuations of the other agents. The center asks the agents for bids. An auction mechanism determines the way the good is allocated, as well as the payment to be made by each agent, based on the information received from the agents. An auction mechanism defines a game. In this game each agent decides on the bid to be made as a function of its valuation of the good. The utility of each agent from the auction is determined by the allocation of the good and by its payment. Most of the literature deals with the case of risk-neutral agents, who have quasi-linear utility functions. In this case, the utility of an agent is v - p if it receives the good, and -p otherwise, where v is its valuation for the good and p is its payment. Economic theory predicts that the agents will behave in an auction according to an equilibrium[1] of the corresponding game. For example, assume that the center sells the good in a second-price auction. In a second-price auction the good is sold to the agent who has offered the highest bid[2], in the second-highest bid. It can be easily checked that truth-revealing is an equilibrium of this auction; it is not worthwhile for an agent to send a bid that differs from its valuation for the good, assuming the other agents submit their actual valuations as their bids. In fact, in this particular case, telling the truth is a dominant strategy; it is not worthwhile to deviate from truth-telling regardless of the behavior of the other agents [Vickrey, 1961].

The objective of the seller is to find an auction mechanism that maximizes its revenue. As shown in [Myerson, 1981], a second-price auction with an appropriate reservation price will do the job (assuming the agents are risk-neutral). The fact that the center obtains the desired outcome, using a mechanism which corresponds to a game where truth-revealing is in equilibrium, is not an accident. The famous revelation principle [Kreps, 1990; Fudenberg and Tirole, 1991] states just that: every objective of the center that is obtained by some mechanism (i.e. in equilibrium of the game associated with that mechanism), can also be obtained by a mechanism in which truth-revealing is an equilibrium of the corresponding game.

12 Optimizing Social Welfare

A second representative mechanism design setting refers to the optimization of social welfare. The optimization of social welfare is a typical objective in resource allocation setups.

Assume that the center needs to select an alternative from a given set of alternatives, *A.* This decision can refer for example to the ordering of jobs in a queue [MacKie-Mason and Varian, 1995; Wilson, 1989], to the selection of the next state to be visited in a search [Ephrati and Rosenschein, 1996], etc. As before, each agent has a certain valuation for each alternative (e.g. an agent may prefer its job to be the first in a queue rather than the second). The social welfare of an alternative is defined as the sum of the agents' valuations for it. The typical center's objective is to select an alternative that maximizes social welfare. Therefore, we are interested in finding a mechanism where each agent reports its valuation, based on which the center can select an alternative that maximizes social welfare. We also require that the mechanism will be budget balanced, i.e. it will not require the center to transfer payments to the agents.

Fortunately, the famous Clarke mechanism [Clarke, 1971] achieves the desired outcome. In a Clarke mechanism each agent is asked to report its valuation, and an alternative a* that maximizes social welfare is selected. Agent *i* is required to transfer a payment of $c_i - d_i$ to the center. This payment is calculated as follows, c_i is the maximal social welfare that can be obtained, when we ignore agent t's valuation; d_i is the sum of the valuations of the other agents for *a8.* It is easy to check that, in the resulting game, truth-revealing (i.e. reporting the true valuation) is a dominant strategy, and that the mechanism optimizes social welfare and is budget balanced.

Both of the above setups are fundamental ones. They capture the idea that a mechanism defines an interaction between the center and the agents, which allows the center to obtain some desired outcome, given that the agents are rational. The revelation principle which we have mentioned before, allows us to restrict our attention to truth-revealing mechanisms, i.e. mechanisms where agents truthfully report their private information. In a typical mechanism, the center will make use of carefully designed payments, in order to obtain the desired behavior in equilibrium of the corresponding game.

3 Mechanism Design Revisited

The mechanism design literature is based on classical economic assumptions. For example, it is built on the assumption

---

[1] We won't present a formal definition of equilibrium here. Roughly speaking, a strategy profile of the agents is in equilibrium, it it is irrational for each agent to deviate from its strategy in that profile, assuming the other agents stick to their strategies.

[2] Coin dipping is used for tie-breaking.

that agents are risk-averse, i.e. have a concave utility function for money. Moreover, as we mentioned, most of the mechanism design literatim deals with risk-neutral agents. Risk neutral agents are an extreme case of risk-averse agents. The Internet serves as an excellent testbed for the study of this and other economic assumptions.

An an example, consider Internet auctions. Auctions are the most widely used (strategic) economic mechanism in the Internet.[3] Massive field experiments in an Internet setup [Lucking-Reiley, 1998] suggest that the classical risk-aversion assumption may be questionable. In particular, consider k-price auctions [Wolfstetter, 1996], In a Ar-price auction a good is sold to the agent with the highest bid, in a price which equals the k-th highest bid. Auction theory predicts that first-price auctions will lead to a higher revenue than second-price auctions when the agents are risk-averse. In addition, it predicts that this difference in revenue becomes more significant when the items that are sold are of higher market value. The above-mentioned field experiments have shown that the latter does not hold, while the former is sometimes refuted as well. This has led some authors to attempt and generate alternative theories of auctions, such as a theory of auctions for risk-seeking agents [Monderer and Tennenholtz, 1998a]. If the agents are risk-seeking, the theory predicts that second-price auctions will be more profitable to the center than first-price auctions. Given that second-price auctions are equivalent, as far as expected revenue is concerned, to classical English auctions[4] (when agents have independent valuations), given that English auctions are the most popular auction mechanism in the Internet, and given that there are many risk-elements in Internet auctions, this alternative theory may deserve further attention. From a design perspective, an important result of this theory is that third-price auctions will yield a higher revenue than the revenue obtained by classical auctions, if the agents are risk-seeking. This point illustrates a way computational implementations may lead us to re-consider the fundamentals of mechanism design.

Computational implementations of economic mechanisms lead to additional basic questions. For example, in the context of auction theory, the following natural questions arise:

1. What is the optimal (revenue maximizing) auction, assuming we have arbitrary risk-averse agents? Do we have to consider new mechanisms or do the classical mechanisms offer optimal revenue?
2. In the Internet, different sellers may compete by selecting different mechanisms for selling similar goods. How is auction theory, and the theory of mechanism design in general, effected by the introduction of competition among mechanisms?

In (Monderer and Tennenholtz, 1998b] we show an upper bound on the optimal revenue obtained by a seller in any auction with risk-averse agents. We show that a second-price auction approaches this upper bound, if the number of agents is large. This may serve as an answer to the first question. In addition, in [Monderer and Tennenholtz, 1998a] we introduce a model for competition among sellers; the sellers compete by selecting among different auction mechanisms. Somewhat surprisingly, we show that in an equilibrium of the competitive setting each seller will choose the same auction it would have chosen if it were a monopoly.

The points we have mentioned so far do not refer to any technological features that are typical to computational environments, such as the Internet. They refer to basic economic questions that are raised and inspired by the implementation of economic mechanisms in computational environments. The discussion of these questions (and their proposed answers) seem however crucial if we wish to make use of economic mechanisms in such environments.

4 From Mechanisms to Protocols

As we mentioned, at first glance it may seem that mechanism design is a different terminology for protocol design, when it is applied to non-cooperative environments. Although this claim is basically correct, researchers have noticed that in order to design protocols for non-cooperative computational environments, various modifications to the mechanism design theory are needed. In particular, researchers have been concerned with the fact agents are resource bounded [Linial, 1992; Sandholm and Lesser, 1997; Kraus and Wilkenfeld, 1991]. This point will be discussed in Section 6. However, except for the important issue of bounded resources there are several other issues that play a major role in the adaptation of economic mechanisms to computational settings. These issues refer to the fact that agents operate in a (computational) distributed system. So, what we need is a model for dealing with the classical and newer problems in distributed systems within the framework of game theory. Such a model is naturally titled a *distributed game* [Monderer and Tennenholtz, 1999a; 1999b]. Hence, a distributed game is a model for dealing with the most general multi-agent interactions in distributed systems. We now briefly discuss two aspects of distributed games that are essential to the design of protocols for non-cooperative computational environments, and have been ignored by the mechanism design literature.

1. In a computational environment protocols are run by agents/processors that are connected by a *communication network.* However, the theory of mechanism design ignores this aspect, and in fact implicitly assumes a very concrete communication network, where each agent is directly connected to the center. This assumption, which is not realistic in most computational settings, is crucial for the mechanism design literature; as it turns out, its adaptation to general computational settings is nontrivial.

---

[4] See e.g. http://www.onsale.com, http://www.ebay.com, http://www.ubid.com.

[4] An English is an open auction. The auctioneer starts with a low asking price, and then continuously increments it as long as there are agents that are still willing to pay the current price. The agent who remained the last, will get the good for the price he was willing to offer.

2. The theory of mechanism design ignores the actual way the information available to the agents is represented, as well as the (bit) structure of the messages. As it turns out, the design of the representation language takes a significant role in the adaptation of economic mechanisms to computational environments.

Consider a distributed system with four processors on a ring, one of them is associated with the center, while the three others are associated with three different agents. Let us denote the center by 0, and the three other agents by 1,2,3. Assume that 1 and 3 are connected directly to 0, while 2 is con nected to 1 and 3. Hence, 1 and 3 can send messages directly to the center, while messages from 2 to the center pass either through 1 or 3. Assume we wish to implement a second-price auction in this communication network. Assume that the valuations of the agents for the good are selected from a uniform distribution on the integers between $\$0$ to $\$ m = 2^k - 1$ for some fixed $k > 0$. The major difference between this setting and the classical mechanism design setting stems from the fact the agents can not directly send messages to the center. For example, if agent 2 sends a bid of $5 to the center through agent 3, then it is only rational for agent 3 to modify this bid to 0. We are interested in (game/information theoretic) techniques that will enable us to solve this problem. Assume a standard binary representation of the valuations/bids. Agent 2 will send to the center two strings, one is a random bit string y[5] which is sent through agent 1, and the other is the bit-by-bit XOR of y with agent 2's valuation for the good, which is sent through agent 3. The center will be able to re-construct agent 2's bid by XORing the two strings, and then will apply the second-price auction. The key point is that any modification by agent 3 of the string sent through it will not increase its expected payoff, since the re-constructed bid will still be uniformly distributed in the given range. Similarly for modifications by agent 1. When these are put into formal terms, we get an implementation of the desired behavior in an equilibrium of the corresponding game.

The above protocol illustrates the role of the communication structure in the transition from mechanisms to protocols. Notice that the above protocol relies on the fact that the valuations are uniformly distributed. If the density of low valuations is smaller than the density of high valuations then it becomes worthwhile for agent 3 to change the string that passes through it to a random string. In order to handle this issue we modify the (syntactic) representation of valuations. We associate with each valuation several strings, where the number of strings that are associated with each valuation is proportional to the density of this valuation. In this case, a uniform selection among strings coincides with a selection of a valuation according to the original distribution. Now, the previous protocol can be modified; y is chosen to be a random string in the allowed representations, and the other string is taken to be the XOR of y with a randomly selected representation of the agent's valuation. The modified protocol enables to implement the second-price auction in the case of non-uniform distributions on the agents' valuations. It illustrates the importance of the actual representation of the agent's valuation, a topic which has been neglected in the mechanism design literature.

The ability to implement the second-price auction by a protocol in the above-mentioned computational environment can be generalized. In [Monderer and Tennenholtz, 1999b] we show that any objective that can be obtained by an economic mechanism, can be also implemented by a protocol in any 2-connected communication graph.

Techniques such as the ones mentioned above enable to extend the theory of mechanism design to arbitrary communication networks. Applications of game-theoretic ideas to communication networks appear in other contexts as well. In particular, much work has been devoted to the study of congestion in communication networks where different agents decide selfishly about the routes of their messages (see e.g. [Korilis eral., 1997]).

5 Asynchronous vs. Synchronous systems

It is common to classify the game theory literature according to several different criteria. One criterion refers to whether the game is static or dynamic. In a static game agents choose actions, based on which they receive particular payoffs. In a dynamic game, the game has several stages, where an agent may select an action in a given stage based on what it observed in previous stages. Another criterion refers to whether the game has complete or incomplete information about the utility function of other agents. Most work in mechanism design fits into the framework of static games with incomplete information. Fortunately, many multi-stage situations (such as the English and Dutch auctions [Wilson, 1992]) can also be modelled as static games.

Multi-agent interactions in computational settings might be static or dynamic/sequential, but may also be parallel. For example, a user may control several agents that participate in several Internet auctions in parallel. Parallel interactions have not played a role in game theory; assume for example that several static interactions are carried out in parallel, then the analysis of each such interaction is not effected by the existence of other parallel interactions. However, real distributed systems are asynchronous, and therefore parallel interactions will in fact take place in some random order. This suggests a new kind of game-theoretic models [Monderer and Tennenholtz, 1999a] that capture the asynchronous nature of distributed systems. As it turns out, asynchronous parallel interactions result in illuminating phenomena. Assume that several similar interactions are held in a parallel (but asynchronous) fashion in several locations. Assume that each user is represented by a set of agents, one at each location, and that the agents can communicate by broadcasting messages. In the framework of such model it can be shown that a cooperative outcome at each of the interactions is obtained in equilibrium. For example, if the famous Prisoners' dilemma is played at each of the locations, then we get an equilibrium in which the

---

[5] In this paper we use the term random selection, in order to refer to a selection based on a uniform distribution.

agents cooperate in all locations. If we have several parallel auctions *(eg.* the famous FCC auction [McMillan, 1994; Milgrom, 1998] made use of parallel interactions), then collusion among the agents at each of these auctions is obtained in equilibrium. Hence, the transition from a synchronous to an asynchronous model plays a most significant role in economic contexts. Needless to say that this treatment is crucial when adapting economic mechanisms to asynchronous distributed systems, such as the Internet.

6 Resource-Bounded Mechanism Design

The fact that agents are resource bounded plays a significant role in computational contexts. Recent work has dealt with resource-bounded mechanism design. A popular topic of study in this context refers to combinatorial *multi-object auctions* (see e.g. [Krishna and Rosenthal, 1996; Wellman *et al.,* 1998]). In a combinatorial auction several objects are sold to a group of agents, where the agents are allowed to make bids for subsets of the goods. Several authors, both in Operations Research [Rothkhof *et al.,* 1998] and in AI [Fujishimaif *al.,* 1999; Sandholm, 1999] have been concerned with the computation of an allocation of goods that maximizes the auctioneer's revenue based on the agents' (combinatorial) bids. Although this problem is NP-hard in general, the related work explores some tractable cases, and points to some useful heuristics.

From the economics perspective, one can view combinatorial auctions as a special case of selection among alternatives. In this case, we may wish to find an allocation of goods which maximizes social welfare, and may wish to use the Clarke mechanism for this purpose. Moreover, Krishna and Perry show [Krishna and Perry, 1998] that if we restrict ourselves to efficient mechanisms (i.e. ones that optimize social welfare), a Clarke mechanism optimizes the seller's revenue when the buyers are risk-neutral. In addition, the Clarke mechanism guarantees that truth-revealing is a dominant strategy for the agents. Notice that if the agents are required to submit bids in a combinatorial auction, and the seller chooses an allocation to optimize its revenue based on these bids, then truth-revealing is generally not in equilibrium.

Given the above, we may wish to consider the implementation of the Clarke mechanism when the agents are resource bounded. The implementation of a Clarke mechanism requires the solution of $n + 1$ optimization problems of the kind discussed in combinatorial auctions. Each of these problems requires the search for an alternative that maximizes the sum of agents' valuations (n problems refer to sets of $n - 1$ agents, and one problem refers to the set of all agents), which is known to be NP-hard. One approach for handling this issue is by using a heuristic for the solution of such optimization problems (notice that the same heuristic serves all the $n + 1$ optimization problems). The problem however is that when using this heuristic we may no longer have truth-revealing as a dominant strategy, and may lose budget balance. Therefore, we are interested in heuristics for optimizing social welfare, that when embodied in a Clarke mechanism will yield a protocol where truth-reveling is a dominant strategy and which is budget-balanced. Several properties of such heuristics, that lead to truth-revealing and budget-balanced protocol, are discussed in [Kflr-Dahav *et al.,* 1998]. In addition; this went presents some settings where any heuristic that does not satisfy the desired properties can be efficiently transformed to a (better) heuristic that does satisfy them.

Another problem that resource-bounded agents need to face is that the number of alternatives (e.g. possible allocations of goods) may be huge, and therefore the agents will not be able to efficiently communicate their preferences. An approach that attempts to decrease the communication burden in the context of a Clarke mechanism is presented in [Ephrati and Rosenschein, 1996].

7 Conclusion

Electronic commerce deals with economic transactions in computational settings. Electronic commerce is a broad area, and there are many different views of this field. This paper has dealt with a concrete aspect of electronic commerce that lies in the intersection of few basic theories; the perspective we take is that a significant part of a theory of electronic commerce should deal with protocols for non-cooperative computational environments. The theory should take into account that the protocols are to be used by different, self-motivated (human) users on one side, but are to run in a computational distributed system on the other side. This calls for the establishment of new unified theories. There are (at least) three centra] topics that these theories need to address:

1. The theory of mechanism design should be carefully revised and extended to deal with (economic) phenomena observed when implementing economic mechanisms in computational settings.
2. Mechanisms should be carefully transformed into protocols, which are to be used in a distributed system. Issues regarding the asynchronous nature of interaction, network topology, and the syntax of representation are crucial in this regard.
3. The effects of the agents* resource bounds on the design of mechanisms should be studied.

Our view is that none of these topics is more important than the others, and that the study of all of them is essential. It is interesting to note that this leaves the border between work in economics/game theory to work in CS/AI somewhat unclear. Indeed, while previous work in economics and game theory has been concerned with resource bounds [Neyman, 1985; Rubinstein, 1986], most of the work in game theory did not refer to basic economic questions inspired by computational settings; similarly, while interesting game-theoretic ideas have been discussed in CS/AI [Boutilier *et al.,* 1997; Rosenschein and Zlotkin, 1994], most work in this regard does not address typical distributed systems features. In this short paper we presented a (biased) perspective on mechanism and protocol design, and illustrated how the above-mentioned issues can be treated. Our aim was not to present anoverview of related research but to give the flavor of the interdisciplinary work carried out in a new challenging area. We believe that bridging the gap between game-theoretic/economic models and working protocols in distributed systems, may result in a significant scientific and technological impact. The newly generated theories may serve as foundations to a theory of electronic commerce.

References

[Bond and Gasser, 1988] A. H. Bond and L. Gasser. *Readings in Distributed Artificial Intelligence.* Ablex Publishing Corporation, 1988.

[Boutilier et a/., 1997] C. Boutilier, Y. Shoham, and M.R Wellman. Special issue on economic principles of multi-agent systems. *Artificial Intelligence,* 94,1997.

[Clarke, 1971] E. Clarke. Multipart pricing of public goods. *Public Choice,* 18:19-33,1971.

[Ephrati and Rosenschein, 1996] E. Ephrati and J. Rosenschein. Deriving consensus in multi-agent systems. *Artificial Intelligence,* 87,1996.

[Fudenberg and Tirole, 1991] D. Fudenberg and J. Tirole. *Game Theory.* MIT Press, 1991.

[Fujishima *et al.,* 1999] Y. Fujishima, K. Leyton-Brown, and Y. Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *IJCAI-99,* 1999.

[Kfir-Dahav et al., 1998] Noa E. Kfir-Dahav, D. Monderer, and M. Tennenholtz. Mechanism Design for Resource-Bounded Agents. Working paper, 1998.

[Korilis *et ai,* 1997] Y. A. Korilis, A.A. Lazar, and A. Orda. Achieving network optima using stackelberg routing stratgies. *IEEE/ACM Transactions on Networking,* 1997.

[Kraus and Wilkenfeld, 1991] S. Kraus and J. Wilkenfeld. The Function of Time in Cooperative Negotiations. In *Prvc. ofAAAI-9l,* pages 179-184,1991.

[Kreps, 1990] D. Kreps. *A Course in Microeconomic Theory.* Princeton University Press, 1990.

[Krishna and Perry, 1998] V. Krishna and M. Perry. Efficient mechanism design. Working Paper, 1998.

[Krishna and Rosenthal, 1996] V. Krishna and R.W. Rosenthal. Simultaneous auctions with synergies. *Games and Economic Behavior,* 17:1-31,1996.

[Linial, 1992] Nathan Linial. Games Computers Play: Game-Theoretic Aspects of Computing. Technical Report 92-5, CS Department, Hebrew University, Israel, 1992.

[Lucking-Reiley, 1998] D. Lucking-Reiley. Revenue equivalence in auctions for real goods: Evidence from an internet-based market. Working paper, Vanderbilt university, March 1998,1998.

[MacKie-Mason and Varian, 1995] J.K MacKie-Mason and H.R. Varian. Some economics of the internet, the 10th Michigan Public Utility Conference, 1995.

[McMillan, 1994] J, McMillan. Selling spectrum rights. *Journal of Economic Perspectives,* 8:145-162,1994.

[Milgrom, 1987] P.R. Milgrom. Auction theory. In T Bewly, editor, *Advances in Economic Theory: Fifth World Congress.* Cambridge University Press, 1987.

[Milgrom, 1998] P. Milgrom. Putting auction theory to work: The simultaneous ascending auction. Technical Report 98-0002, Department of Economics, Stanford University, 1998.

[Monderer and Tennenholtz, 1998a] D. Monderer and M. Tennenholtz. Internet Auctions. Working paper, 1998.

[Monderer and Tennenholtz, 1998b] D. Monderer and M. Tennenholtz. Optimal Auctions Revisited. In *Proceedings of AAAI-98,* 1998.

[Monderer and Tennenholtz, 1999a] D. Monderer and M. Tennenholtz. Distributed Games, to apear in Games and Economic Behavior, 1999.

[Monderer and Tennenholtz, 1999b] D. Monderer and M. Tennenholtz. Distributed Games: From Mechanisms to Protocols. In *AAAI-99,* 1999.

[Myerson, 1981] R. Myerson. Optimal auction design. *Mathematics of Operations Research,* 6:58-73,1981.

[Neyman, 1985] A. Neyman. Bounded complexity justifies cooperation in the infinitely repeated prisoner's dilemma. *Econ. Lett.,* 19:227-229,1985.

[Rosenschein and Zlotkin, 1994] Jeffrey S. Rosenschein and Gilad Zlotkin. *Rules of Encounter.* MIT Press, 1994.

[Rothkhof *et al.,* 1998] M.H. Rothkhof, A. Pekec, and R.M. Harstad. Computationally manageable combinatorial auctions. *Management Science,* 44(8):1131-1147,1998.

[Rubinstein, 1986] A. Rubinstein. Finite automata play the repeated prisoner's dilemma. *Journal of Economic Theory,* 39:83-96,1986.

[Sandholm and Lesser, 1997] Tuoams W. Sandholm and Victor R. Lesser. Coalitions Among Rationally Bounded Agents. *Artificial Intelligence,* 94(1):99-137,1997.

[Sandholm, 1999] T. Sandholm. An algorithm for optimal winner determination in combinatorial auctions. In *IJCAI-99,*1999.

[Vickrey, 1961] W. Vickrey. Counterspeculations, auctions, and competitive sealed tenders. *Journal of Finance,* 16:15-27,1961.

[Wellman *etal,* 1998] M.P. Wellman, P.R. Wurman, W.E. Walsh, and J.K. MacKie-Mason. Auction protocols for distributed scheduling. Working paper, 1998.

[Wilson, 1989] W. Wilson. Efficient and competitive rationing. *Econometrica,* 57(1): 1-40,1989.

[Wilson, 1992] R. Wilson. Strategic analysis of auctions. In R J. Aumann and S. Hart, editors, *Handbook of Game Theory,* volume 1. Elsevier Science Publisher, 1992.

[Wolfstetter, 1996] E. Wolfstetter. Auctions: An introduction. *Journal of Economic Surveys,* 10(4):367-420,1996.