Games and Economic Behavior **35,** 1–5 (2001)
doi:10.1006/game.2001.0848, available online at http://www.idealibrary.com on IDEAL®

INTRODUCTION

Economics and Artificial Intelligence

Dov Monderer and Moshe Tennenholtz

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

and

Hal Varian

School of Information Management and Systems, 102 South Hall, University of California, Berkeley, California 94720-4600

When this special issue was announced, we had to explain why we thought that the growth of the Internet would strongly influence the relations between computer scientists and economists. By now, such an explanation is hardly necessary.

Within the broad fields of economics and computer science, the subspecialties that have the most points of contact are Game Theory and Artificial Intelligence. We use the term "Artificial Intelligence" for the science and technology of building automated decision-makers that can act effectively in an environment consisting of both human and artificial decision-makers. We use the term "Game Theory" to describe the science that analyzes environments in which the outcome depends on the strategic decisions of a group of decision-makers.

Relations between game theory and computer science have been established in the past. An excellent example of the deep connection between game theory and computer science is illustrated by the TARK—Theoretical Aspects of Rationality and Knowledge—conferences (former title: Theoretical Aspects of Reasoning about Knowledge, e.g., Halpern, 1986). These conferences have emphasized works on the foundations of knowledge and rationality that are of interest to both fields. Such works naturally appear in this issue. Papers on knowledge and rationality in this issue include the following:

0899-8256/01 $35.00

Binmore and Samuelson, in “Coordinated Action in the Electronic Mail Game,” analyze a version of the coordinated attack problem (Gray, 1978; Yemini and Cohen, 1979; Halpern and Moses, 1990) that is known in economics as the “electronic mail game” (Rubinstein, 1989). See also Monderer and Samet (1989). In this version the noisy communication technology is voluntary and costly. The authors show, in particular, the existence of equilibrium in which only one message is sent. They present an evolutionary stability condition (Selten, 1980) that eliminates the equilibrium in which all messages are ignored, and analyze other equilibrium profiles that survive the stability criterion.

Heifetz and Mongin, in “Probability Logic for Type Spaces,” use a logical approach (see, e.g., Kripke, 1963; Aumann, 1995) to knowledge theory to give an axiom system that is sound and complete with respect to the class of type spaces in the sense of Harsanyi (1967–1968). They compare their work to the approach to type space taken in computer science (Fagin and Halpern, 1994).

Lehmann, in “Expected Qualitative Utility Maximization,” presents a model for rational decision making—Expected Qualitative Utility Maximization—that encompasses the Maximin criterion and generalizes expected utility maximization. The properties of this model are studied in the framework of von Neumann and Morgenstern (1947). Its main ingredient is the definition of a qualitative order on nonstandard models of real numbers and the consideration of nonstandard utilities.

However, the emergence of the Internet, as well as recent work carried out by computer scientists, AI researchers, and economists, has shown that the spectrum of mutual interest is much broader. In particular, it includes the field of learning in multiagent systems (see, e.g., Fudenberg and Levine, 1998, for a survey of learning in games from the economics perspective, and Kaelbling *et al.*, 1996; Shavlik and Dietterich, 1990, for partial surveys on machine learning in AI). Papers on learning in this issue include the following:

Amy Greenwald, Eric J. Friedman, and Scott Shenker, in “Learning in Network Contexts: Experimental Results from Simulations,” describe the results of simulations of a suite of learning algorithms in network contexts. They explore the extent to which the asymptotic play depends on three factors: limited information, asynchronous play, and degree of responsiveness of the learning algorithm.

Parkes and Huberman, in “Multiagent Cooperative Search for Portfolio Selection,” present a new multiagent model for the multiperiod portfolio selection problem; in such a model we have a system of bounded-rational cooperative agents that pool their initial wealth, manage a share of that investment each, and then pool their final wealth. The authors present a quantitative assessment of the performance of their model in simulated stock markets.

The third part of this special issue consists of six papers on mechanism design in a computational setting. This is a new field in Artificial Intelligence that builds itself on the theory developed in economics and both extends and applies it in various ways. This field had been established before this special issue (see, e.g., Varian, 1995) and there is considerable ongoing work in this area (see, e.g., Tennenholtz, 1999). The papers on mechanism design in this issue are as follows:

Nisan and Ronen, in "Algorithmic Mechanism Design," introduce new ways to deal with the classical scheduling problem using the theory of mechanism design in economics; they also introduce two ideas new to this field: a measure of optimality and a definition of probabilistic mechanism.

Shoham and Tennenholtz, in "On Rational Computability and Communication Complexity," deal with the questions of what a market can compute and at what expense. The authors treat rationality as a computational resource and show that it has interesting relationships with other computational resources, in particular with communication complexity. The authors also describe mechanisms of low communication complexity for information elicitation.

Sandholm and Lesser, in "Leveled Commitment Contracts and Strategic Breach," define and analyze the equilibrium of a new type of contract that can be implemented in automatic negotiations. They define a leveled commitment contracting mechanism that allows agents to capitalize on uncertain future events by having the possibility of unilaterally decommitting from a contract based on local reasoning.

Wellman, Walsh, Wurman, and MacKie-Mason, in "Auction Protocols for Decentralized Scheduling," deal with the problem of efficient scheduling of time slots. As the "optimal" mechanism for this problem—the generalized Vickrey mechanism—is not computationally feasible, the authors deal with another mechanism—an ascending auction. They use the best-reply dynamic to analyze the outcome of such auctions. In addition, because every such auction results in a vector of prices, one for each good, the authors deal with the general theory of discrete markets that are cleared by prices.

Wurman, Wellman, and Walsh, in "A Parametrization of the Auction Design Space," deal with practical problems in implementing auctions. They present an extensive breakdown of the auction design space that captures the essential similarities and differences of many auction mechanisms. This parametrization serves as an organizational framework in which to classify work within the field and uncovers new, potentially useful, mechanisms.

Vulkan, in “Equilibria in Automated Interactions,” describes software agents that play games. The author shows that such agents can get more than humans because they are able to prove to each other that they do not use their information. This “power” of ignorance resembles that discussed in Monderer and Tennenholtz (1999), who describe software agents that can commit themselves not to remember and are therefore able to cooperate in repeated prisoner dilemma-type games.

Other relevant issues that belong to the field of economics and CS are not covered here, for example, the topic of “playing theoretically simple but computationally hard” games (e.g., chess). This topic has been extensively studied in AI (see, e.g., Russell and Norvig, 1995, for a general discussion of game-playing in the context of standard AI techniques), but only recently studied by game theorists (e.g., Jehiel and Samet, 2000).

This special issue illustrates some of the existing and emerging connections between computer science and Artificial Intelligence particularly in regard to economics and game theory. These connections build on a shared perspective: dealing with decision-makers who need to make decisions in a potentially uncertain environment and may learn about their environment and interact with one another. Fifty years ago, several of the founders of game theory (e.g., John Nash and Lloyd Shapley) and Artificial Intelligence (e.g., John McCarthy and Marvin Minsky) were developing their early theories as colleagues in the same research institute—Princeton University. During the subsequent years, these fields went their own ways, refining these early achievements. Now, 50 years later, it seems that close relationships between the fields are again emerging. It is our belief that the level of this interaction will continue to increase, leading to significant scientific and technological contributions.

REFERENCES

Aumann, R. (1976). “Agreeing to Disagree,” *Ann. Statist.* **4**, 1236–1239.

Aumann, R. J. (1995). “Interactive Epistemology,” Discussion Paper 67, Center for Rationality and Interactive Decision Theory, The Hebrew University.

Fagin, R., and Halpern, J. Y. (1994). “Reasoning about Knowledge and Probability,” *J. ACM* **41**(2), 340–367.

Fudenberg, D., and Levine, D. K. (1998). *The Theory of Learning in Games*. Cambridge, MA: MIT Press.

Gray, J. (1978). “Notes on Database Operating Systems,” in *Operating System: An Advanced Course* (R. Bayer, R. M. Graham, and Seegmuller, Eds.), New York: Springer-Verlag.

Halpern, J. Y. (1986). “Theoretical Aspects of Reasoning about Knowledge,” *Proceedings of the 1986 Conference*. San Francisco: Morgan Kaufmann.

Halpern, J. Y., and Moses, Y. (1990). “Knowledge and Common Knowledge in a Distributed Environment,” *J. ACM* **37**, 549–587.

Harsanyi, J. C. (1967–1968). "Games with Incomplete Information Played by Bayesian Players," Parts I–III, *Management Sci*. **14**, 159–182, 320–334, 486–502.

Jehiel, P., and Samet, D. (2000). *Learning to Play Games in Extensive Form by Valuation*. Tel Aviv: Tel Aviv University.

Kaelbling, L. P., Littman, M. L., and Moore, A. W. (1996). "Reinforcement Learning: A Survey," *J. AI Res*. **4**, 237–285.

Kripke, S. A. (1963). "Semantical Analysis of Modal Logic. I. Normal Modal Propositional Calculi," *Z. Math. Logik Grundlagen Math*. **9**, 67–96.

Monderer, D., and Samet, D. (1989). "Approximating Common Knowledge with Common Beliefs," *Games Econom. Behav*. **1**.

Monderer, D., and Tennenholtz, M. (1999). "Distributed Games," *Games Econom. Behav*. **28**, 55–72.

Rubinstein, A. (1989). "The Electronic Mail Game, Strategic Behavior under Almost Common Knowledge," *Am. Econom. Rev*. **79**, 385–391.

Russell, S., and Norvig, P. (1995). *Artificial Intelligence: A Modern Approach*. New York: Prentice Hall.

Selten, R. (1980). "A Note on Evolutionary Stable Strategies in Asymmetric Animal Contests," *J. Theoret. Biol*. **84**, 93–101.

Shavlik, J. W., and Dietterich, T. G. (1990). *Readings in Machine Learning*. San Francisco: Morgan Kaufmann.

Tennenholtz, M. (1999). "Electronic Commerce: From Game-Theoretic and Economic Models to Working Protocols," in *Proceedings of the National Conference on Artificial Intelligence* (*AAAI99*).

Varian, H. (1995). "Economic Mechanism Design for Computerized Agents," in *Proceedings of Usenix Conference on Electronic Commerce*, *New York*, *July 11–12, 1995*.

von Neumann, J., and Morgenstern, O. (1947). *Theory of Games and Economic Behavior*, 2nd ed. Princeton Univ. Press.

Yemini, Y., and Cohen, D. (1979). "Some Issues in Distributed Processes Communication," in *Proceedings of the 1st International Conference on Distributed Computing Systems*, pp. 199–203.