

# Agents and Markets

Amy Greenwald, *Brown University*

Nicholas R. Jennings, *University of Southampton*

Peter Stone, *University of Texas at Austin*



**A**utonomous agents are intelligent software programs. Typically, agents are situated in an environment. Repeatedly, they sense their environment, engage in some decision making whereby they select actions, and execute their actions, which, in turn, impact their environment. Moreover, in most cases, the environment contains a number of such agents whose actions typically mutually affect one another. This interdependence arises because the different agents, with their own aims and objectives, must operate in a common environment that has finite resources and capabilities. Depending on the dependency's nature, several different types of social interaction occur between the agents, including

- *Cooperation* (working together to achieve a common objective)
- *Coordination* (arranging the actions to ensure they're coherent)
- *Negotiation* (coming to a mutually acceptable agreement on some matter)

Against this background, one of the most important mechanisms for exchanging resources or conducting negotiations among multiple individuals is the marketplace. Examples include open-air bazaars in which buyers and sellers haggle informally over the prices of trinkets; Sotheby's Auction House, which uses English outcry auctions to sell high-priced goods to select buyers; and the New York Stock Exchange, where traders exchange millions of shares every day. Traditionally, the individuals participating in markets were assumed to be (boundedly) rational human beings acting in their own self-interest. Increasingly, however, autonomous agents are becoming active participants in marketplaces. Given this trend, this special issue devotes itself to studies on the interactions between autonomous agents and markets.

### Autonomous agents vs. human traders

Market agents have at least three main potential advantages over human traders. First, agents have computational advantages. They can operate faster and handle more transactions. However, and perhaps more important, in many trading contexts agents can more effectively optimize the set of goods to buy or sell given current prices. For example, in a travel scenario, which involves combinatorially many packages comprising flights, hotel rooms, rental cars, and so on, an autonomous agent should be able to search through the alternatives and make the most advantageous trades more effectively than humans. (One of the articles in this issue considers just such a scenario.)

Second, agents don't get distracted—they can participate in ongoing auctions without directing their attention elsewhere. Take, for example, bidding agents on eBay. Such auctions last for several days with new bids arriving anytime. If a human bidder happens to be offline when an auction closes and if other bidders place critical bids at that time, the first bidder, who has no opportunity to react, possibly forgoes the opportunity to improve upon his or her bid and so derive some benefit (or utility). However, you can program

an autonomous bidding agent to top any new bid up to some maximum price, thus ensuring that the bidder wins the auction and that the bidder's utility is maximized, as long as the price is right.

Third, you can program an agent to be immune to the flaws of reasoning to which humans are notoriously susceptible. Daniel Kahneman and Amos Tversky demonstrated that people can make different decisions regarding two choices with the same outcome.<sup>1</sup> For example, most people would rather take a sure \$500 than a 50-percent chance at \$1,000. On the other hand, if people are initially given \$1,000, they would rather take a 50-percent chance at losing \$1,000 than a sure \$500 loss. Even though from a rational perspective, these two decisions should be the same, people

One of the most important mechanisms for exchanging resources or conducting negotiations among multiple individuals is the marketplace.

nonetheless decide differently. Agents, on the other hand, can be programmed to avoid this problem.

Rajarshi Das and his colleagues have illustrated several of these advantages in the trading domain.<sup>2</sup> They pitted humans and agents against each other in controlled experiments in a continuous double-auction scenario (like the stock market). They discovered that a significant amount of human-agent trading occurred (not just agents trading with agents and humans trading with humans) and that the agents performed significantly better, even when the humans were experienced.

Despite autonomous agents' potential advantages, many challenges hinder their widespread deployment in marketplaces. For example, agents must be reliable and trustworthy. If a human bidder is relying on an agent to place bids in an auction, it's unacceptable for the agent to crash. It's also unacceptable for the agent to reveal any informa-

tion about its future bids to competing bidders, lest its competitors exploit this information to their advantage. Furthermore, we need to devise methods of accurately encoding preference relations in agents that will represent human bidders.

This special issue contains articles that specifically address these and other outstanding challenges that we must overcome before agents become a ubiquitous component of electronic marketplaces.

### This issue

The articles in this special issue cover a broad range of topics. In "Computational-Mechanism Design: A Call to Arms," Rajdeep K. Dash, David C. Parkes, and Nicholas R. Jennings present an overview of the terminology and issues in designing the rules that govern interactions between agents (this is the problem of *computational mechanism design*). They outline ways that mechanisms can and should be tailored to address computational concerns, and they advocate the design of mechanisms suited to distributed systems.

From an architectural perspective, before agents can engage in a market, they must know where to find it. For example, if you deploy an agent as a shopbot to buy a particular book, it first must know where to find the online bookstores. For this purpose, *middle agents* can direct buyers toward potential sellers and vice versa. In "The Role of Middle-Agents in Electronic Commerce," Itai Yarom, Claudia V. Goldman, and Jeffrey S. Rosenschein explain the role of various types of middle agents in e-commerce settings. They analyze, through simulation, the effect of various configurations of buyers, sellers, and intermediaries.

The remaining four articles examine specific examples of agents participating in marketplaces.

In "Protocols for Negotiating Complex Contracts," Mark Klein and his colleagues propose a protocol for negotiations between two agents in which several related concerns exist. For example, an item's or service's price is typically related to its quality (higher quality means higher price), and delivery time is often related to price (buyers must often pay a premium to have a service carried out immediately). The authors propose a protocol that relies on a trustworthy mediator that searches through the solution space to find an agreement that's nearly Pareto optimal. This scheme aims to

avoid outcomes like that which arises in the prisoners' dilemma. And perhaps what's more interesting, the authors also describe an unmediated negotiation protocol that performs as well as the mediated negotiation protocol.

In "Trading Agents Competing: Performance, Progress, and Market Effectiveness," Michael P. Wellman and his colleagues analyze the third international Trading Agent Competition. TAC is intended to pose benchmark problems in the autonomous-agents research community. Started in 2000, the classic TAC scenario has simulated travel agents bidding in simultaneous auctions to obtain flights, hotel rooms, and entertainment tickets for clients with varying preferences. Strategies have evolved over the years, the success of which depends heavily on an agent's ability to predict auction clearing prices. But up to this point, there's been little or no analysis of the relative performance of different prediction strategies. In addition, there's been no analysis of the agents' overall efficiency as they collectively solve TAC's distributed-allocation problem. This article provides both types of analyses.

In a similar vein is a new project at the University of Pennsylvania that introduces another platform for agent competitions, but in a stock-trading scenario with market conditions that are directly tied to NASDAQ in real time. Stock ticker data has been readily available for several years, but only recently has real-time information about complete standing-order books (that is, buy and sell limit orders) been obtainable via electronic crossing networks such as Island. The Penn-Lehman Automated Trading Project takes advantage of this recent advance to create a market simulation in which (simulated) trading agents can buy and sell stock, both from one another and from true traders. The stock prices in the simulation are thus tied directly to the real world. In "The Penn-Lehman Automated Trading Project," Michael Kearns and Luis Ortiz introduce this novel simulator, presenting its underlying implementation details and relating some of the early experiences of autonomous-agent participants.

In recent years, some countries have deregulated their electricity supply industries to obtain the benefits of increased competition. This has given birth to the application area of competitive electricity markets. In "MASCHEM: A Multiagent That Simulates

Competitive Electricity Markets," Isabel Praça and her colleagues describe a simulator that can serve as a decision support tool. With this tool, users can explore both the effects of various rules that could be imposed on the marketplace and the effects of different bidding strategies on the overall system's performance.

**W**e believe that the coupling of autonomous agents and marketplaces will fundamentally change how many goods and services are traded. In particular, we believe dynamic pricing will once again become the norm for many types of trading encounters (fixed prices are, in fact, a relatively new phenomena) and that trading relationships will be established and disbanded in a much more agile and timely fashion than is currently possible. This vision relies on markets' intrinsic economic properties and on having autonomous participants that can rapidly make informed decisions in ever-changing environments.

To make this vision a reality, much work remains. Such work needs to address the scientific aspects associated with making decisions in complex situations, the engineering aspects of building reliable software in challenging environments, and the social aspects of trusting software to make (financial) decisions on behalf of users. These various strands of endeavor are well represented by the articles in this special issue, which, we believe, provides an important point of departure for work in this exciting new area. ■

**References**

1. D. Kahneman and A. Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, vol. 47, no. 2, Mar. 1979, pp. 263-291.
2. R. Das et al., "Agent-Human Interactions in the Continuous Double Auction," *Proc. 17th Int'l Joint Conf. Artificial Intelligence (IJCAI 01)*, Morgan Kaufmann, 2001, pp. 1169-1187.

The Authors



**Amy Greenwald** is an assistant professor of computer science at Brown University. Her primary research area is the study of economic interactions among computational agents.

Her primary methodologies are game-theoretic analysis and simulation. Her research is applicable in areas ranging from dynamic pricing via pricebots to autonomous bidding agents. She recently received an NSF career grant entitled "Computational Social Choice Theory" and was named a Computing Research Association Digital Government Fellow. She previously researched information economies at IBM's T.J. Watson Research Center. Her paper "Shopbots and Pricebots" (joint work with Jeff Kephart) was named Best Paper at IBM Research in 2000. Contact her at the Computer Science Dept., Brown Univ., Box 1910, Providence, RI 02912; amygreen@cs.brown.edu.



**Nicholas R. Jennings** is a professor of computer science in Southampton University's School of Electronics and Computer Science, where he carries out basic and applied research in agent-based

computing. He is also the Deputy Head of School (Research); the head of the Intelligence, Agents, Multimedia Group; and the chief scientific officer for Lost Wax. He received his PhD in artificial intelligence from the University of London. He won the Computers and Thought Award in 1999, an IEE Achievement Medal in 2000, and the ACM Autonomous Agents Research Award in 2003. Contact him at the School of Electronics and Computer Science, Univ. of Southampton, Southampton SO17 1BJ, UK; nrj@ecs.soton.ac.uk.



**Peter Stone** is an assistant professor in the Department of Computer Sciences at the University of Texas at Austin, where he investigates multiagent learning. His research interests include planning,

machine learning, multiagent systems, robotics, and e-commerce. He received his PhD in computer science from Carnegie Mellon University. Contact him at the Dept. of Computer Sciences, Univ. of Texas at Austin, Austin, TX 78712-1188; pstone@cs.utexas.edu.