

Software agents and the route to the information economy

Jeffrey O. Kephart*

Institute for Advanced Commerce, IBM Research, P.O. Box 704, Yorktown Heights, NY 10598

Humans are on the verge of losing their status as the sole economic species on the planet. In private laboratories and in the Internet laboratory, researchers and developers are creating a variety of autonomous economically motivated software agents endowed with algorithms for maximizing profit or utility. Many economic software agents will function as miniature businesses, purchasing information inputs from other agents, combining and refining them into information goods and services, and selling them to humans or other agents. Their mutual interactions will form the *information economy*: a complex economic web of information goods and services that will adapt to the ever-changing needs of people and agents. The information economy will be the largest multiagent system ever conceived and an integral part of the world's economy. I discuss a possible route toward this vision, beginning with present-day Internet trends suggesting that agents will charge one another for information goods and services. Then, to establish that agents can be competent price setters, I describe some laboratory experiments pitting software bidding agents against human bidders. The agents' superior performance suggests they will be used on a broad scale, which in turn suggests that interactions among agents will become frequent and significant. How will this affect macroscopic economic behavior? I describe some interesting phenomena that my colleagues and I have observed in simulations of large populations of automated buyers and sellers, such as price war cycles. I conclude by discussing fundamental scientific challenges that remain to be addressed as we journey toward the information economy.

My colleagues[†] and I believe that, over the course of the next decade or two, the world economy and the Internet will merge into an *information economy* bustling with billions of autonomous software agents that exchange information goods and services with humans and other agents. Software agents will represent—and be—consumers, producers, and intermediaries. They will facilitate all facets of electronic commerce, including shopping, advertising, negotiation, payment, delivery, and marketing and sales analysis.

In the information economy, the plenitude and low cost of up-to-date information will enable consumers (both human and agent) to be better informed about products and prices. Likewise, producers will be better informed about and more responsive to their customers' needs. Low communication costs will greatly diminish the importance of physical distance between trading partners. These and other reductions in economic friction will be exploited and contributed to by software agents that will respond to new opportunities orders of magnitude faster than humans could.

The agents that we envision will not be mere adjuncts to business processes. They will be economic software agents: independent self-motivated economic players, endowed with algorithms for maximizing utility and profit on behalf of their human owners. From other agents, they will purchase inputs, such as network bandwidth, processing power, or database access rights, as well as more refined information products and

services. They will further refine and add value to these inputs by synthesizing, filtering, translating, mining, or otherwise processing them and will sell the resultant product or service to other agents. In essence, these agents will function as automated businesses that create and sell value to other agents and businesses and in so doing will form complex efficient economic webs of information goods and services that adapt responsively to the ever-changing needs of humans for physical and information-based products and services. With the emergence of the information economy will come previously undreamt-of synergies and business opportunities, such as the growth of entirely new types of information goods and services that cater exclusively to agents. Ultimately, the information economy will be an integral and perhaps dominant portion of the world's economy.

How might this vision of the future be realized? In this paper, I discuss a plausible route toward such a future, using a chronological series of examples that begin with present-day Internet trends and leading successively toward a full-fledged information economy. I begin with the example of two Internet startups that revolve around eBay as a basis for discussing some technological, business, and legal trends that hint at a future in which economic software agents will become important. Then, I use the results of recent laboratory experiments in which we pitted automated bidding agents against human bidders to illustrate how and why economic software agents are likely to supplant human economic decision makers. This demonstration forms the basis for a further extrapolation, in which economic software agents will be so numerous that interactions among them will have a significant effect on the economy. I illustrate some of the generic modes of collective behavior that one might observe in markets dominated by economic software agents, using a simulation study of a market in which the vendors all use automated price-setting algorithms. I conclude with a discussion of some significant scientific challenges that lie ahead.

Trends

As a basis for discussing various technological, business, and legal trends that point toward an eventual information economy, I shall focus on examples involving two Internet startups that revolve around eBay, the Internet's dominant business-to-consumer auction site.

This paper results from the Arthur M. Sackler Colloquium of the National Academy of Sciences, "Adaptive Agents, Intelligence, and Emergent Human Organization: Capturing Complexity through Agent-Based Modeling," held October 4–6, 2001, at the Arnold and Mabel Beckman Center of the National Academies of Science and Engineering in Irvine, CA.

Abbreviations: CDA, continuous double auction; GD, Gjerstad and Dickhaut; ZIP, Zero Intelligence Plus; MY, myopic best-response algorithm; DF, derivative follower; NIR, no internal regret.

*E-mail: kephart@us.ibm.com.

[†]This paper represents the collective work and thought of many colleagues at IBM Research, the University of Michigan, and Brown University, whose names appear in the acknowledgments section.

eSnipe is an Internet business that automates a common bidding strategy on eBay called sniping. Sniping is the practice of waiting until a few seconds before the close of an auction before submitting one's bid to prevent others from countering. In a fascinating empirical and theoretical study (available at <http://www.economics.harvard.edu/~aroth/papers/ebay-revisedOriginal.pdf>), Roth and Ockenfels have found that sniping is surprisingly common and successful (1). Rather than submitting their bids directly to eBay, eSnipe subscribers indicate a bid amount and a buffer time t , and eSnipe automatically places the bid at a time t before the auction's close. Currently, eSnipe reports 50,000 subscribers, 10,000 bids submitted per day, and a total of \$200 million dollars in bids placed through it per year. eSnipe plans to extend its service to cover auction sites other than eBay in the near future.

Originally, eSnipe's automated bidding service was offered to bidders for free, and the site was supported by revenue from online advertisements. However, under new ownership,[‡] eSnipe has begun to charge for winning bids at a rate of 1% of the bid amount for winning bids over \$25, up to a maximum fee of \$10.

eSnipe exemplifies two distinct aspects of economic software agents that will be important in the information economy. First, it hints at the coming importance of agent-mediated electronic commerce, in which software agents will automate or assist in making and acting on economic decisions. Although eSnipe merely carries out the direct explicit instructions supplied to it by human owners, it clearly points the way toward more sophisticated bidding agents endowed with adaptive strategies that will participate in a variety of different auctions, both in the business-to-consumer and business-to-business realms. Buy.com is an early sign of related developments in automated posted pricing: it uses software that implements "we will not be undersold" by scanning competitors' prices and undercutting them by a small margin. Revenue management systems in use by the airline industry for years clearly demonstrate the level of sophistication that can be attained and the value that can be gained by using automated posted pricing. It is easy to imagine that such techniques could be applied much more generally to a wide variety of industries.

Second, eSnipe is an early example of an information service that charges for what it provides, rather than using the typical ad-based Internet business model. I believe that charging for information services is an essential development, and that the successful emergence of the information economy depends on the widespread adoption of this practice.

To understand why this is so, consider the following entry in eSnipe's Frequently Asked Questions:

"Q. You should show the current price of the item on eBay so I can know whether I need to increase my bidding amount."

"A. Including the current price in the My Bids section sounds good, but it poses some surprising problems. The most important is that our servers would have to check eBay's servers hundreds of thousands of times a day, with no benefit to eBay."

"eSnipe prefers that you use eBay as much as possible, so we include the link to each item, allowing you to check the current price quickly and easily. It's better for eBay (you're likely to see their ads and announcements that way), [and] better for eSnipe (we don't anger them too much by polling their servers constantly)."

eSnipe is clearly concerned about how eBay would react if eSnipe were to poll its site too frequently, and with good reason. In 1999, eBay barred several auction aggregator sites (which function like search engines for auctions), including Auction-Watch, RubyLane, and BiddersEdge, from searching its site. All had been guilty of a breach of ordinary "netiquette:" they had

ignored the "robot exclusion headers" (2) that eBay had inserted into their web pages to warn away web crawlers. BiddersEdge ignored eBay's warning, and in December 1999, eBay sued BiddersEdge for unlawful trespass, charging that they were automatically polling eBay about 100,000 times per day in response to queries from their customers.

In May 2000, Judge Ronald M. Whyte (3) issued an injunction preventing BiddersEdge from scanning eBay's pages directly. Judge Whyte stated that, in his opinion, whereas the extra load placed on eBay by BiddersEdge was not significantly harmful in itself,

"If BE's (BiddersEdge) activity is allowed to continue unchecked, it would encourage other auction aggregators to engage in similar recursive searching of the eBay system, such that eBay would suffer irreparable harm from reduced system performance, system unavailability, or data losses."

Eugene Volokh (3), a law professor at University of California, Los Angeles, commented that

"One direct consequence of this opinion is that it shows that commercial sites will be able to block this kind of intelligent agent access, spider access, by comparison services. That means that some of these comparison sites will go out of business, or they may have to pay money to get permission to use spiders and abide by various restrictions."

The last clause of Volokh's comment supports my contention that information services (like eBay) should charge other information services (like BiddersEdge) for the information they provide. Although information resources might be provided to people for free in hopes that the investment will be repaid by the selling of banner ads, this business model fails with agents, which tend not to read ads! If agents representing information services pay one another for the information services that they use, then the right incentives will be in place, and our envisioned economic web of information goods and services can flourish (4). Information and network resources will not be abused, or made unavailable, as was the case with eBay and BiddersEdge.[§] Furthermore, human owners of information services would undoubtedly welcome increased agent-based demand for their resources, rather than trying to ward it off with exclusion headers or lawsuits.

Thus we anticipate that, regardless of whether its core business is an e-commerce function (like eSnipe's bidding service), a search engine, an application service provider, or anything else, an information service provider will charge for the service it provides and pay for the services it uses. Automated businesses of all sorts will automatically buy and sell information services, and therefore they will be economic agents. Will we be able to endow them with a sufficient amount of economic intelligence to do the job? The next section suggests that we have every reason to be optimistic.

The First Economic Software Agents

For the information economy vision to be realized, agents must attain a level of economic performance that rivals or exceeds that of typical humans; otherwise, people would not entrust agents with making economic decisions on their behalf. In this section, I describe a series of experiments that demonstrate the superiority of software agents to humans in a commonplace price-setting domain: the continuous double auction (CDA).

The CDA is an interesting choice for several reasons. First, it is used extensively in the real world, particularly in financial markets such as NASDAQ and the New York Stock Exchange.

[‡]Ultimately, the breakdown in cooperation ended up being even more severe. BiddersEdge and eBay settled out of court in March 2001, whereupon BiddersEdge promptly left the auction aggregation business altogether, depriving many customers of a worthwhile service and depriving eBay of the extra business that was brought to them via BiddersEdge.

[‡]Fittingly, the new owner bought eSnipe on eBay, using eSnipe to submit the winning bid!

Second, there is an extensive literature on both all-human experiments [dating back to Vernon Smith's original work in 1962 (5)] and all-agent experiments [dating back at least as far as the Santa Fe Institute Double Auction Tournament in the early 1990s (6)], and ranging up to the present-day Trading Agent Competition conducted by Michael Wellman and his team at the University of Michigan (7, 8).

The CDA is a double-sided market: buyers place bids indicating the amount they are willing to pay for an item, and sellers place "asks" indicating the amount for which they would be willing to sell an item. When a bid and an "ask" match or overlap, a trade is executed immediately, hence the "continuous" in the term CDA.

Experimental Setup. Our experiments took place at the Watson Experimental Economics Laboratory at IBM Research in Yorktown Heights. The lab is specially designed for a variety of economics experiments involving human subjects. We modified the existing software to accommodate a set of software bidding agents that we had designed. Each experiment involved six humans and six agents participating in a multiunit CDA, with half of each type playing as buyers and the other half playing as sellers. Some of the human subjects were students from local community colleges; others were IBM researchers. Following standard experimental economics practice, human participants were paid in proportion to their surplus, i.e., their gains from trade. This experimental protocol motivates human subjects to play as if the limit prices assigned by the experimenter really reflected their underlying value (if a buyer) or cost (if a seller).

A typical experiment consisted of 9–16 3-min periods. At the beginning of each period, each participant was assigned a set of limit prices, one for each of several identical units that one might buy (or sell). A seller's limit prices typically increased for each successive unit sold, reflecting increasing production costs. A buyer's limit prices typically decreased with each successive unit bought, reflecting demand saturation. Each human buyer or seller had a unique set of limit prices that was fixed for the duration of a period but was changed randomly every few periods. To enable fair comparison between agent and human performance, each human's set of limit prices was mirrored by one of the agents throughout the experiment. The figure of merit for each participant was the total surplus they earned.

Each human buyer or seller was provided with a graphical user interface (GUI) that supplied information about limit prices, the current state of the market (the bid/ask queue), and historical price trends in the market. The GUI also provided means for entering bids or "asks." Agent buyers and sellers were provided with exactly the same set of information that was used to generate the GUI for human participants; this information was fed directly to the agents, who assimilated it and acted on it asynchronously, in real time.

Agent Strategies. Agent strategies were based on two CDA strategies that had been published previously: the Zero Intelligence Plus (ZIP) strategy introduced by Cliff and Bruten (9) and the GD strategy introduced by Gjerstad and Dickhaut (10). We made numerous changes and enhancements to both strategies to take into account different market conditions (11).

For example, ZIP was originally developed for a call market, in which trades are executed only at the end of the period rather than continuously. Therefore, the time at which a bid was made was not an issue—only the price was important. In a continuously clearing market like the CDA, bid timing is very important, and it can be tricky. We introduced "fast" and "slow" variants of both ZIP and GD. The fast variants considered placing bids as soon as they were informed of new market activity. Interestingly, fast bid timing had the potential to create avalanches of bidding, as agents responded to one another's responses. The slow

variants reconsidered their options at a rate based on an internal timer—typically every 5 sec.

Another significant set of modifications dealt with the fact that both ZIP and GD were originally designed for markets without a persistent order queue. In other words, either an order traded, or it was supplanted by a better one. However, for our experiments, we used a more realistic version of the CDA with a persistent order queue: it retains orders that either are or used to be the best bid or "ask." Use of the different variant of the CDA necessitated several modifications to both ZIP and GD. For example, GD maintains statistics of prices at which bids have been accepted or rejected. Unmatched bids in an order queue are in an ambiguous state—neither accepted nor rejected—introducing subtleties and pathologies into the GD method that required modifications (11).

Experimental Results. To establish a baseline, we first ran all-agent experiments with homogeneous populations of ZIP or GD agents. We found robust approximate convergence to the theoretical equilibrium prices, with the agents obtaining efficiencies of between 0.98 and 1.00. (The efficiency is the ratio of the total surplus to that which would have been attained had all trades been executed at the theoretical equilibrium price.) This value can be compared with prior measurements of the efficiency of all-human populations, which tend to be slightly lower, in the range of 0.95–0.98.

Then we ran six experiments as described above, with various selections of agent types. Fig. 1 shows a typical example of the time series of trade prices obtained from one experiment with six fast-GD agents and six humans.

Our main observations about the six experiments are as follows:

- The average surplus earned by agents was 3–38% higher than that obtained by humans, with the difference averaging about 20%.
- In the latter phases of the experiments, this discrepancy was reduced to about 5–7%. However, further experiments with experienced humans failed to show any additional improvement.
- Agents tend to fare considerably better than their human counterparts on an individual basis as well. In a few experiments, one of the six humans was able to just edge out his or her agent counterpart, but typically the agents beat their human counterparts substantially.
- The proportion of trades occurring between agents and humans was substantial, typically in the range of 35–45% of all trades, as compared with the 50% that would be obtained if there were no bias. This finding demonstrates that agent and human markets are strongly coupled, despite the fact that agents can respond much more quickly to market events.
- Trades that are far removed from the theoretical equilibrium price are mostly agent traders taking advantage of poor decisions by human traders, although the magnitude of these blunders tends to decrease as humans learn over the course of the experiment.
- The greater the proportion of agents, the closer the market comes to perfect efficiency, resulting (ironically) in better performance by humans. This effect was first observed in an experiment in which only three humans participated and achieved an overall efficiency of just 3% less than the agents. Subsequent informal experiments have borne out this effect.
- An algorithm that fares well against other reasonably strong agent opponents will not necessarily do well against humans. In the slow-ZIP experiments, highly nonoptimal human strategies fooled the ZIP algorithm into behaving even more poorly than humans on the sell side of the market. We were able to diagnose the problem and modify ZIP to avoid such pitfalls,




Fig. 1. Trade price vs. time for experiment of October 25 with six GD fast agents and six humans. The vertical dashed lines indicate the start of a new trading period. The 16 trading periods were divided into four phases, each with its own set of limit prices. In each phase, p^* is shown by the horizontal dashed lines. Trades between two agents are shown with open circles, between two humans with open squares, and between an agent and a human with solid circles. The labels AB, AS, HB, and HS refer to the average efficiency of Agent Buyers, Agent Sellers, Human Buyers, and Human Sellers, respectively. The overall market efficiency was 0.946:1.052 for the agents and 0.840 for the humans.

but the point remains valid: in testing agent algorithms, there is no substitute for doing real experiments with humans.

Prognosis. The agents we have developed are certainly not the best we can envision. We have already begun to develop a stronger bidding algorithm, called GDX, that resoundingly beats GD and ZIP. Informal tests indicate that GDX fares even better against humans than GD. Over time, we and the community as a whole will develop stronger bidding strategies for the CDA and a variety of other auction institutions. Thus human performance can be expected to lag even further behind that of agents! Consequently, we believe that agents will, to a large extent, supplant humans in CDA and other auctions. Extending these conclusions one step further, there is every reason to believe that economic software agents will ultimately prove superior to humans in other economic domains as well, including posted pricing and negotiation.

This is not to say that humans will have no role in economic decision making. We believe humans and agents will combine their strengths. Market participants will really be human-agent hybrids, with agents executing low-level strategies at a fast time scale and humans executing higher-level metastrategies that involve setting the parameters of agent bidding strategies or selecting new agent bidding strategies. Perhaps agent-human hybrid strategies will continue to coevolve against one another in a perpetual arms race.

Emergent Behavior of Economic Software Agents

As economic software agents continue to grow in sophistication, variety, and number, interactions among them will start to become significant. These interactions will be both direct and indirect. Direct interactions among agents will be supported by a number of efforts that are already under way, including standardization of agent communication languages, protocols, and infrastructures by organizations such as the Foundation for Physical Agents (FIPA) (<http://www.fipa.org>) and Object Management Group (OMG) (<http://www.omg.org>), myriad attempts to establish standard ontologies for numerous products and markets (CommerceNet being one prominent player in this arena; see <http://www.commerce.net>), and the development of various micropayment schemes such as IBM MicroPayments (12). Indirect interactions among agents will occur through the medium of the economy itself. For example, two sellers competing for market share may find that their actions are strongly

coupled, even if they do not communicate with one another directly.

When interactions among agents become sufficiently rich, some crucial qualitative changes will occur. One effect is that new classes of agents will be designed specifically to serve the needs of other agents. Among these will be “middle agents” (13) that supply brokering and matchmaking services to other agents for a fee.

Another effect, to which the rest of this section is devoted, is the appearance of interesting collective modes of economic behavior that arise from direct or indirect interactions among economic software agents. Economic software agents will differ in significant respects from their human counterparts, particularly in the speed with which they can sense and respond to changes and new opportunities. These differences are likely to translate into new forms of collective economic behavior that are not typically observed in today’s human-dominated economy.

To help anticipate emergent phenomena that may be commonplace in markets dominated by software agents, we have analyzed and simulated several different market models. Several phenomena appear to be quite generic: not only are they observed in many different scenarios, but their root causes appear to be very basic and general. The remainder of this section summarizes two such phenomena that are particularly interesting and ubiquitous: price war cycles, and an effect that we call “the prisoner’s metadilemma.”

Price War Cycles. We have studied several market models in which many buyers purchase commodities or multiattribute goods from two or more sellers. In these models, buyers are assisted by automated purchasing algorithms that continually monitor prices and product attributes and purchase an item from the vendor that best satisfies the individual buyer’s utility function. Several distinct variants of this model have been explored: a market for a simple commodity (in which some buyers perform only a limited search among the set of sellers) (14, 15); a vertically differentiated market (in which the good may be offered at different levels of quality by different sellers) (16); and information filtering (17, 18) and information bundling models (19) (in which the seller sets both the price and the parameters that determine the product’s configuration).

We have experimented with several different pricing algorithms in each of these models. One generic pricing algorithm that applies in all of these scenarios is the “myoptimal” or myopic

best-response algorithm (MY). The MY algorithm assumes perfect knowledge of the aggregate buyer demand, from which it computes the profit landscape. The profit landscape is the profit as a function of the price vector consisting of all of the sellers' prices (or the price and product parameter vector, if the good is multiattribute and configurable). The MY algorithm is conceptually simple: it selects the price (or price/product parameter set) that yields the highest profit, given the current offerings of the other sellers.

If all sellers use MY, then the same basic pattern emerges in all of the models that we have explored. In markets in which the sellers set the price only, the sellers' prices cycle indefinitely in a sawtooth pattern. Initially, all sellers charge the price that would be set by a monopolist; this is followed by a phase in which the price drops linearly as the sellers undercut one another in an effort to grab most of the market share. Finally, at a price somewhat above the marginal production cost for the good, the sellers stop undercutting one another and suddenly raise their price back to the monopolist level, whereon the cycle begins anew. This sudden jump in price occurs at the point where the margin is so low that grabbing most of the market share is no longer the best policy; instead, it is best to make a large profit on the relatively small market segment consisting of a portion of buyers who do not diligently search for the best deal.

In models in which sellers can set product configuration parameters (such as the subject matter of news or journal articles in either the information filtering or information bundling models), we observe an analog of this cyclical pattern that occurs in the higher-dimensional space of price and product parameters. Typically, price is finely quantized, and product parameters are discrete-valued. In such a case, a typical price trajectory consists of linear drops punctuated by sharp discontinuities up or down. These discontinuities coincide with quick shifts by all sellers to a new set of product parameters. This discontinuous pattern will eventually repeat, although the number of discontinuous phases within one cycle may be quite large. The consequences for buyers in such markets could be dire: rather than covering a broad range of niches, the sellers are, at any given moment, focused on one particular product configuration, leaving many buyers unsatisfied.

Although the intuitive explanations for the sudden price or price/product discontinuities differ considerably across the various models, there is a generic mathematical principle that explains why this behavior is so ubiquitous and provides insight into how broadly it might occur in real agent markets of the future. Mathematically, the phenomenon occurs in situations where the underlying profit landscape contains multiple peaks. Competition among sellers collectively drives the price vector in such a way that each seller is forced down a peak in its profit landscape, although its individual motion carries it up the peak. At some point, a seller will find it best to abandon its current peak and make a discontinuous jump to another peak in its landscape. This discontinuous shift in the price vector suddenly places all other sellers at a different point in their own profit landscapes. If the next seller to move finds itself near a new peak in its landscape, it will make an incremental shift in its price and hence in the price vector; otherwise, it will respond with yet another radical shift.

Although the models we have explored are highly idealized, we believe that price or price/product cycles have the potential to occur in approximate form in real agent markets of the future. The insights we have obtained from our theoretical studies suggest that markets are susceptible to cyclical price/product wars when the following conditions are all met:

- A profit landscape exists, i.e., there is a predictable relationship between the price vector and each sellers' profit. For this condition to be satisfied, buyers must be extremely responsive

to changes in the price vector, i.e., there must be no inertia. The existence of automated buying agents may move the economy much closer to this ideal than it is at present.

- The profit landscape contains multiple peaks. This condition is virtually assured in any realistic situation, because such peaks tend to be generated easily. Consider for example the common preference for goods that are cheaper and suppose there are two sellers. Then Seller 1's profit landscape has a huge discontinuity at $p_1 = p_2$: profits are much higher on the $p_1 < p_2$ side of the $p_1 = p_2$ line than they are on the $p_1 > p_2$ side.
- The sellers' pricing algorithm permits a global search. One option is for sellers to either know or learn the profit landscape well enough to perform at least an approximate global optimization of it. This option gives them the capability of jumping discontinuously to a (temporarily) better solution. Another option demands less knowledge and ability on the part of the sellers. Even if they have little understanding of the global features of the landscape, all they really need is an adventurous spirit that causes them to at least occasionally experiment with radically different prices or price/product parameters. When they "get lucky," this will cause the same dramatic shifts, although the classic sawtooth pattern will be replaced by an irregular pattern of jump discontinuities with random jitters (representing exploration) superimposed on it.

Prisoner's Metadilemma. Another general principle that emerges from our studies is that a strategy (say for pricing) that is superior when used in isolation may fare poorly if that strategy becomes widely adopted, possibly leading to an overall decline in social welfare.

A specific scenario in which we have observed this phenomenon is in the commodity market model mentioned in the previous subsection. As was stated, universal adoption of the MY strategy leads to cyclical price wars. We have also studied a different strategy, called "derivative follower," or DF, that makes less stringent assumptions about the knowledge available to the agent. DF does not rely on information about buyer demand or even other sellers' prices. It simply continues to move the price in one direction as long as the profit increases, reversing direction when the profit decreases. When all sellers use DF, the resulting emergent behavior is quite interesting, with lots of randomized upward and downward trends in price. On average, however, the prices (and therefore the profits) are sustained at a higher level than is achieved by a population consisting entirely of MY sellers.

Suppose that a single MY seller is introduced into a previously all-DF market. Then MY tends to just undercut the lowest DF seller. Because the DF sellers tend to maintain fairly high prices, the MY seller's average profit is much higher than that of the DF sellers. This profit differential creates an incentive for a DF seller to shift to the MY strategy. If all sellers are free to choose their pricing strategies, they will opt for MY, even if others have already chosen MY. This incentive to shift strategies would drive the market to a state in which all sellers use MY, hence earning lower profits than they would have obtained had they all made a pact to use DF.[†]

This phenomenon is strongly reminiscent of the prisoner's dilemma, a well known non-zero-sum game in which rational players end up with a lower payoff than would be obtained by a pair of irrational players (31). Interestingly, however, it occurs at the metalevel: choosing a pricing strategy rather than a specific price. For this reason, we refer to it as the "prisoner's metadilemma."

The prisoner's metadilemma has been observed in other

[†]Note that such a pact would not be collusive in the usual sense, as the sellers are agreeing only on their pricing strategies, not on the prices themselves.

situations as well. Suppose that the MY strategy is not feasible because buyer demand cannot be known or discovered explicitly. Then one can use a machine learning method called “no internal regret” (NIR) that learns a probabilistic pricing strategy (20). An interesting form of emergent behavior actually makes NIR a better individual player against DF than MY. DF responds to NIR by setting its prices at the monopolist level, permitting NIR to choose a price probability distribution with a preponderance of mass just below the monopolist price. Thus there is a strong incentive for DF sellers to switch to NIR. Even when this is done asynchronously, regardless of the mix of NIR and DF players in the market, there is at every step of the way an incentive for a player to switch from DF to NIR, until ultimately all sellers play NIR. At this point, the average profit sinks to a level significantly lower than when all sellers use the DF (or the MY) strategy.

A similar phenomenon has been observed in an auction scenario in which automated bidders use a variety of strategies. Sophisticated strategies can fare well when used by only a few agents but can lead to lower social welfare when overused (21, 22). It has also been observed in a multiagent resource allocation scenario (23).

New Vistas and Scientific Challenges

The information economy will be by far the largest multiagent system ever envisioned, with numbers of agents running into the billions. Economic software agents will differ from economic human agents in significant ways, and their collective behavior may not closely resemble that of humans. It would be imprudent to use the world’s economy as an experimental testbed for software agents. Analysis and simulation are valuable tools that can be used to anticipate some of the opportunities and problems that lie ahead and to gain insights that can be incorporated into the design of economic agents and mechanisms.

Our various attempts to model the future information economy suggest some fundamental scientific challenges that need to be addressed. Several of these have to do with understanding and harnessing collective effects that will become important after economic software agents become sufficiently abundant that their mutual interactions will affect markets and, ultimately, the entire world economy. I shall conclude this paper by mentioning a few of these challenges, which can also be viewed in a more positive light as new vistas for scientific exploration.

The ability to create a new nonhuman breed of economic player may engender a whole new engineering discipline within economics. Rather than regarding mathematical models of economic agents as approximate descriptions of human behavior, they can be taken more seriously as precise prescriptions of economic software agent behavior. It is likely that many existing economic models and theories will be dusted off and extended in new directions that are dictated by the new demands that will be placed on them. They will be taken as blueprints, not approximations. An example that we have pursued in our work is the theory of price dispersion, which has traditionally been an academic exercise aimed at explaining how prices in a competitive market can remain unequal and above the marginal production cost. Traditionally, it was natural to assume that search costs were fixed and determined exogenously. In contrast, we have used price dispersion theory as a launching point for exploring how an information agent could strategically set search costs to maximize its profits, a line of inquiry that has led to the discovery of several interesting nonlinear dynamical effects (24).

Another serious set of challenges and new opportunities for fundamental scientific developments arises in the realms of machine learning and optimization. Economic software agents will need to use a broad range of machine learning and optimization techniques to be successful. They will need to learn, adapt to, and anticipate changes in their environment. Many of these environmental changes will result directly or indirectly from

adaptation by other agents that are similarly using machine learning and optimization. Yet much of the work on machine learning and optimization (and practically all of the theorems) assumes a fixed environment or opponent.

One example of economic machine learning and optimization that we have explored pits two adaptive price setters against one another (25, 26). Rather than using the MY strategy, which chooses a price that optimizes immediate profits, the agents use a reinforcement learning technique called “Q-learning” to maximize anticipated future discounted profits (32). Ordinary single-agent Q-learning is guaranteed to converge to optimality. Very little is known about the behavior of Q-learning in a multiagent context. We observed that two competing Q-learning agents converged to a fixed pricing strategy in some restricted situations (the prices themselves followed a price war cycle of greatly reduced amplitude). However, the agents were much more likely to exhibit cycling in the pricing strategies (not the prices themselves); that is, they altered the way they responded to one another’s prices episodically in an approximately cyclical pattern. Interestingly, the two agents settled into different patterns, despite the absence of any inherent asymmetry, and the time-averaged profits of the agents were improved by the Q-learning technique despite the lack of convergence. In general, understanding the dynamic interactions among a society of learners is of fundamental theoretical and practical interest, and only a few beginning efforts have been made in this area (25, 27–29).

A second example of the complex interplay between learning, optimization, and dynamics that can occur among interacting economic software agents was explored in an information bundling context (30). An agent representing an information bundling service expressed its pricing policy as a set of parameters that dictated a nonlinear dependency on the type and number of articles offered in each category. The agent’s objective was to maximize profits. To do so, it experimented with various settings of price parameters, observed the resulting profit, and used the well-known amoeba nonlinear optimization technique to try to learn and climb the profit landscape. Simultaneously, the buyers attempted to learn their utility functions for the offered product by sampling it. Statistical fluctuations would sometimes lead buyers to falsely conclude that the prices were too high to justify a purchase. Once the buyers stopped purchasing articles, they had no further source of information and would not reenter the market unless the prices were lowered. In effect, the profit landscape was changing continually, and the amoeba algorithm’s implicit assumption that it was optimizing a static function caused it to perform extremely poorly. Although the buyers’ actual valuations were in fact static, their own learning and experimentation created a dynamic that flummoxed the optimization algorithm, resulting in pricing behavior that exacerbated the problem rather than correcting it. Although we were able to correct this specific problem by making the amoeba algorithm periodically reevaluate its old samples, it seems likely that entirely new approaches involving a unified understanding and treatment of learning, optimization, and dynamics will have to be invented.

Fundamental scientific work in these and other areas will have enormous implications for the future of electronic commerce and the entire world economy.

This paper describes the collaborative efforts of several researchers in the Information Economics group at the IBM Thomas J. Watson Research Center and in the Computer Science and Economics Departments and the School of Information at the University of Michigan. My colleagues and coauthors at IBM have included: Amy Greenwald (now at Brown University), Jim Hanson, Gerry Tesaro, Rajarshi Das, David Levine, Richard Segal, Benjamin Grosz (now at the Massachusetts Institute of Technology Sloan School), and Jukka

Sairamesh. Special thanks go to Steve White, who helped conceive the Information Economies project and has remained a guiding influence throughout its existence. Finally, I thank IBM's Institute for Advanced

Commerce, Steve Lavenberg, and the rest of the Adventurous Research Program Committee for several years of financial support and guidance.

1. Roth, A. & Ockenfels, A. (2002) *Am. Econ. Rev.*, in press.
2. Eichmann, D. (1995) *Computer Networks and ISDN Systems* **28**, 127–136.
3. Judge Bans Web Site's Use of eBay Data (*New York Times*), May 26, 2000, C7.
4. Kephart, J. O. (2000) in *Cooperative Information Agents IV, Lecture Notes in Artificial Intelligence*, eds. Klusch, M. & Kerschberg, L. (Springer, Berlin), Vol. 186.
5. Smith, V. L. (1962) *J. Polit. Econ.* **70**, 111–137.
6. Rust, J., Miller, J. & Palmer, R. (1994) *J. Econ. Dyn. Control* **18**, 61–96.
7. Wellman, M. P., Wurman, P. R., O'Malley, K., Banera, R., de Lin, S., Reeves, D. & Walsh, W. E. (2001) *IEEE Internet Computing* **5**, 43–51.
8. Greenwald, A. & Stone, P. (2001) *IEEE Internet Computing* **5**, 52–60.
9. Cliff, D. & Bruten, J. (1997) *Minimal-Intelligence Agents for Bargaining Behaviors in Market-Based Environments. Technical Report HPL-97-91* (Hewlett Packard Laboratories, Bristol, U.K.).
10. Gjerstad, S. & Dickhaut, J. (1998) *Games Econ. Behav.* **22**, 1–29.
11. Das, R., Hanson, J. E., Kephart, J. O. & Tesauero, G. J. (2001) in *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, ed. Nebel, B. (Morgan Kaufmann, San Francisco), pp. 1169–1176.
12. Herzberg, A. & Yochai, H. (1997) *Proceedings of the Sixth International World Wide Web Conference*, eds. Genesereth, M. R. & Patterson, A. (Elsevier, Amsterdam).
13. Decker, K., Sycara, K. & Williamson, M. (1997) in *Proceedings of the 15th International Joint Conference on Artificial Intelligence (IJCAI-97)* (Morgan Kaufmann, San Francisco).
14. Greenwald, A. & Kephart, J. O. (1999) in *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99)* (Morgan Kaufmann, San Francisco), pp. 506–511.
15. Greenwald, A. R., Kephart, J. O. & Tesauero, G. T. (1999) in *Proceedings of First ACM Conference on Electronic Commerce*, ed. Wellman, M. (ACM, New York), pp. 58–67.
16. Sairamesh, J. & Kephart, J. O. (1998) in *Proceedings of First International Conference on Information and Computation Economies (ICE-98)* (ACM, New York).
17. Kephart, J. O., Hanson, J. E. & Sairamesh, J. (1998) *Artificial Life* **4**, 1–23.
18. Kephart, J. O., Hanson, J. E., Levine, D. W., Grosz, B. N., Sairamesh, J., Segal, R. B. & White, S. R. (1998) in *Proceedings of the Second International Workshop on Cooperative Information Agents (CIA '98)*, eds. Klusch, M. & Weiss, G. (Springer, Berlin), pp. 160–171.
19. Kephart, J. O. & Fay, S. (2000) in *Proceedings of the Second ACM Conference on Electronic Commerce* (ACM, New York), pp. 117–127.
20. Greenwald, A. R. & Kephart, J. O. (2001) in *Proceedings of the Fifth International Conference on Autonomous Agents* (ACM, New York), pp. 560–567.
21. Bye, A., Preist, C. & Bartolini, C. (2001) in *Proceedings of the Fifth International Conference on Autonomous Agents* (ACM, New York), pp. 545–551.
22. Park, S. Durfee, E. H. & Birmingham, W. (1999) in *Proceedings of the Third International Conference on Autonomous Agents* (ACM, New York).
23. Kephart, J. O., Hogg, T. & Huberman, B. A. (1990) *Physica D* **42**, 48–65.
24. Kephart, J. O. & Greenwald, A. R. (2002) in *Autonomous Agents and Multi-Agent Systems: Special Issue on Game-Theoretic and Decision-Theoretic Agents*, in press.
25. Tesauero, G. J. & Kephart, J. O. (2002) in *Autonomous Agents and Multi-Agent Systems: Special Issue on Game-Theoretic and Decision-Theoretic Agents*, in press.
26. Kephart, J. O. & Tesauero, G. J. (2000) in *Proceedings of the Seventeenth International Conference on Machine Learning (ICML'2000)* (Morgan Kaufmann, San Francisco), pp. 463–470.
27. Vidal, J. M. & Durfee, E. H. (1998) in *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS '98)* (IEEE Computer Society Press, Los Alamitos, CA), pp. 317–324.
28. Hu, J. & Wellman, M. P. (1998) in *Proceedings of the Second International Conference on Autonomous Agents (Agents '98)* (ACM, New York), pp. 239–246.
29. Tesauero, G. J. & Kephart, J. O. (1998) *Proceedings of First International Conference on Information and Computation Economies (ICE-98)* (ACM, New York).
30. Kephart, J. O., Das, R. & MacKie-Mason, J. K. (2000) *Agent-Mediated Electronic Commerce, Lecture Notes in Artificial Intelligence* (Springer, Berlin).
31. Axelrod, R. (1984) *The Evolution of Cooperation* (Basic Books, New York).
32. Watkins, C. J. C. H. & Dayan, P. (1992) *Machine Learn.* **8**, 279–292.