An Agent-Based Global Economy

Michael N. Huhns • *University of South Carolina* • *huhns@sc.edu*

Many people are buying music CDs over the Web, but with multiple retail sites offering the same CDs, how do you decide which retailer to buy from? Typically, you visit a few sites and choose the one with the lowest price. *Shopbots*—shopping agents that automatically search the Internet to obtain information about prices and other attributes of goods and services—are ideal helpers for such a task.[1] The better ones can visit hundreds of sites, giving price-conscious consumers a powerful tool that could work to the detriment of some retailers.

However, retailers also use shopbots (they call them *pricebots*) to check each other's prices. This is precisely how Books.com used to operate: it would charge one cent less than the cheapest price offered by Amazon, Barnes & Noble, and Borders. With this strategy, a clever retailer such as Books.com might garner sufficient sales to capture a disproportionate share of the market.

Visions of the Future

But what happens if all shoppers use shopbots and all retailers use pricebots? This is the subject of current research by Jeff Kephart and his colleagues at IBM,[2,3] who build simulated worlds of shoppers and retailers and then unleash teams of agents into those worlds. Depending on the conditions they impose, price wars erupt, chaos ensues, and some retailers go out of business. In many ways, these simulations preview the future of electronic commerce.

The Internet is moving toward an open, friction-free marketplace in which software agents will manage the buying and selling of goods. In the near term, agents' major roles will occur in the general economy of goods and services, where they will dynamically link teams of buyers and sellers, producers and consumers, in efficiently managed transactions. Agents will contribute to this economy by reducing communication and interaction costs—the "friction"—associated with doing business.

In the new *information economy*, agents will purchase ideas, computational results, search results, and raw information from other agents and refine this information through aggregation, filtering, knowledge discovery, data mining, classification, translation, and other processes into information products. They will then sell the resultant products to other agents—both computational and human.

Earlier Agents on the Web columns[4,5] reported on how agents are managing Web-based auctions such as eBay and how teams of agents are using collaborative filtering to recommend books at Amazon and Barnes & Noble. Some years hence the Internet will likely support numerous agents that may not be associated with a particular website. These independent, economically motivated software agents will find and process information and disseminate it to humans and, increasingly, to other agents.

Using the Internet for an agent-based information economy seems beneficial. After all, economic processes already manage the conflicting needs of billions of human agents. Adam Smith referred to just such processes when he explained in *The Wealth of Nations* (1776), "every individual necessarily labors to render the annual revenue of the society as great as he can. . . He intends only his own security; and by directing that industry in such a manner as its produce may be of the greatest value, he intends only his own gain, and he is in this, as in many other cases, led by an invisible hand to promote an end which was no part of his intention." Perhaps a similar process will operate in an agent-based economy.

Simulating Agent-Based Economies

University of Michigan professor Michael Wellman has turned Smith's idea into a computational mechanism, albeit on a limited scale, in his development of market-based programming.[6] He set up a system of producing and consuming agents, each intent on maximizing its own profit. By running the system until it reaches equilibrium, he can solve optimization problems.

At first glance, it seems that this mechanism can be applied to agent-based information economies. However, information economies involve human as well as software agents, and by comparison, software agents are immeasurably less sophisticated, less flexible, less able to learn, and notoriously lacking in common sense. Given these differences, it's entirely possible that agent-based economies will behave in unexpected ways.

This is exactly what Kephart and his colleagues found.[3] They first assume a mixture of buyers, some searching for the lowest price and the rest buying from the first seller who meets the

price they have in mind. Their price-bots operate according to one of three different pricing strategies:

- *game-theoretic*, based on a mixed-strategy Nash equilibrium;
- *myopically optimal*, by which prices are set according to the known mix of buyers and other sellers' current prices; and
- *derivative*, whereby each seller sets prices according to its own profit trend—increasing its prices as long as profits are rising and decreasing prices as long as profits are falling.

The first two strategies require perfect knowledge of buyer statistics and the number of sellers, whereas the third requires only local knowledge.

When all pricebots follow the same strategy, the derivative strategy yields the highest profits for all. However, with a mixture of pricebots following different strategies, those following the myopically optimal strategy fare best. Interestingly, if pricebots are allowed to change their strategies, they will end up with the myopically optimal strategy, which leads not only to repeated price wars but also to lower profits for all.

The Future of Brand Names

One potential outcome of a mass market of interacting pricebots and shopbots is the loss of brand-name identity.[7] In our present economy, sellers rely heavily on brand-name influence, and buyers learn to depend on known brands. First-generation bots focus only on price and are oblivious to brand names: to them, a mom-and-pop garage operation is indistinguishable from a major manufacturer, and there is no such thing as brand loyalty.

A March 2000 survey by the Boston Consulting Group found that convenience was the foremost motivator for Internet purchasing. First-generation bots are meeting this need, but consumers have other concerns. Buying at the lowest price isn't very appealing if your goods never arrive, their quality is poor, or they can't be returned if they prove unsuitable. Under these circumstances, dealing with a reputable seller can be more important than simply getting the lowest price. Next-generation bots should be able to help with this aspect of e-commerce as well.

One way they might accomplish this is through online certification services, much as sellers now rely on consumer credit ratings when extending credit. Early examples of certification services are BBBOnLine (http://www.bbbonline.org), provided by the Council of Better Business Bureaus, and TRUSTe (http://www.truste.org). These services certify the privacy and/or fiscal responsibility of retail sites, but not the quality of their service. When no certifying organization can be found for sellers of a desired product or service, information can be obtained from a site such as Dun & Bradstreet (http://www.dnb.com/), which provides financial reports on companies and organizations, both domestic and international. By incorporating certification services, agents may be able to retain brand-name value.

Matching Buyers and Sellers

Building on general ideas pioneered by Gnutella, FreeNet, and SETI@home, Mojo Nation (http://www.mojonation.net) is creating a digital marketplace for the exchange of idle disk space, bandwidth, and CPU cycles. By providing services and resources to other Mojo Nation peers, participants earn credit they can use to consume content or sell for cash. Although right now the marketplace is limited to mostly computational goods, it provides a strategy for micropayments among participants and a new currency: the Mojo.

An alternative to using or inventing a currency is a return to the types of markets that originated during the Middle Ages, when trade was based on bartering—the direct exchange of goods with no intermediate currency ("I'll trade my extra pig for your extra goat"). Because an agent-based economy provides direct links between producers and consumers, without retailers and other middlemen, it might also significantly reduce the need for traditional currencies. This is impractical in our current markets, because while individuals and corporations simultaneously produce and consume, they typically produce narrowly and consume broadly. Pairwise matches are thus unlikely. I might like to trade my extra chickens to Ford for a new car, but Ford needs steel, not chickens. Agents, however, might be able to dynamically form large coalitions and chains of producers and consumers, then find matches to make such bartering feasible.

Agents will continue to evolve from assistants into decision-makers, and their autonomy and responsibility will increase. Ultimately, transactions among economic software agents will constitute an essential and perhaps even dominant portion of the world economy. ∎

REFERENCES

1. B. Krulwich, "The BargainFinder Agent: Comparison Price Shopping on the Internet," in *Agents, Bots, and Other Internet Beasties*, J. Williams, ed., Sams.Net Publishing, Macmillan, Indianapolis, Ind., 1996, pp. 257-263.
2. A.R. Greenwald and J.O. Kephart, "Shopbots and Pricebots," *Proc. 16th Int'l Joint Conf. Artificial Intelligence*, Morgan Kaufmann, San Francisco, Calif., 1999, pp. 506-511.
3. J.O. Kephart, J.E. Hanson, and A.R. Greenwald, "Dynamic Pricing by Software Agents," to appear in *Computer Networks*, 2000; available online at http://www.research.ibm.com/infoecon/paps/html/rudin/rudin.html.
4. M.N. Huhns and A.K. Malhotra, "Negotiating for Goods and Services," *IEEE Internet Computing*, Vol. 3, No. 4, July/Aug. 1999, pp. 97-99.
5. M.N. Huhns and J.M. Vidal, "Online Auctions," *IEEE Internet Computing*, Vol. 3, No. 3, May/June 1999, pp. 103-105.
6. M. Wellman, "Market-Based Programming," in *Readings in Agents*, M.N. Huhns and M.P. Singh, eds., Morgan Kaufmann, San Francisco, Calif., 1997.
7. C. Bayers, "Capitalist Econstruction," *Wired*, Vol. 1, No. 6, Mar. 2000; available online at http://www.wired.com/wired/archive/8.03/markets.html.

Michael N. Huhns is professor of computer science and engineering at the University of South Carolina, where he also directs the Center for Information Technology.