

Game Theory and the Design of Electronic Markets

Amy Greenwald

Brown University

1. Introduction

In electronic commerce (e-commerce), the negotiation of an exchange of goods or services is conducted electronically, rather than physically. Today, the most prevalent form of electronic market is the electronic auction house. eBay, whose consumer auction infrastructure is facilitating consumer-to-consumer transactions, earned revenues of \$1.3 billion in 2005. Google and Yahoo!, whose electronic-sponsored search auctions are facilitating business-to-consumer and business-to-business transactions, earned revenues of \$6.1 and \$5.3 billion, respectively, in 2005.

Game theory is the study of decision making among multiple agents (human or software). In a game, the decisions of all agents jointly determine the outcome. An auction is a classic example of a game: the bids of all participants determine the winner and the winning price. This essay illustrates some of the features and bugs in the design of electronic consumer auctions (eBay's and Amazon's) and sponsored search auctions (Yahoo!'s, and Google's) by applying game-theoretic reasoning. The Trading Agent Competition (TAC), a game-based platform for joint experimentation with intelligent software agent strategies and electronic market designs, is also described. Simulation environments such as TAC are indispensable because the effectiveness of an electronic market design is contingent on the behavior of the market's participants.

By respecting game-theoretic principles and by analyzing electronic market designs using agent-based simulations, markets can be made more efficient so that, in the future, they do a better job of matching buyers with sellers than what is achieved today.

2. Consumer Auctions

eBay. A standard mechanism for selling an item on eBay is via an ascending auction that terminates at a fixed time, with the highest bidder winning the item at the winning bid price. eBay auctions are monitored by autonomous bidding agents called proxy bidders. After a user enters his willingness to pay (WTP) for an item, a proxy bidder bids on the user's behalf, as long as the WTP exceeds the winning bid.

To make use of proxy bidding, eBay provides the following guidelines, which instruct bidders to bid their true WTPs for items up front:

Enter the highest amount you'd be willing to pay for the item. eBay will automatically raise your bid, only as much as is needed for you to remain the high bidder. (This is called proxy bidding...) This means that you may win an item for less than your maximum bid.

However, it is not optimal to follow eBay's guidelines [1].

Imagine Alice and Bob are the only two participants in an auction. Suppose Alice's WTP A exceeds Bob's WTP B . Alice is a manual bidder: she monitors the auction herself, and manually enters a bid one increment above the

Related Resources:

Web Pages

[Trading Agent Competition \(TAC\)](#): the home page for TAC, including news, specifications, and competition information.

[Al Roth's Game Theory, Experimental Economics & Market Design Page](#): contains links to papers and introductory information about game theory, experimental economics, and market design.

Articles

[Computational-mechanism design: a call to arms](#) Dash, R.K., Jennings, N.R., Parkes, D.C. *IEEE Intelligent Systems*

[Position auctions](#), Varian, H. Technical paper, UC Berkeley

Books

[Combinatorial auctions](#) Cramton P., Shoham Y., Steinberg R. (Eds.), 2006

[Handbook of computational economics, vol. 2: agent-based computational economics](#) Tesfatsion L., Judd K.L. (Eds.), 2006

Conferences and Workshops

[ACM Conference on Electronic Commerce \(ACM EC\)](#): an annual conference sponsored by the ACM Special Interest Group on Electronic Commerce (SIGECOM)

[Trading Agent Design and Analysis Workshop \(TADA\)/Agent Mediated Electronic Commerce Workshop \(AMEC\)](#): a

highest bid p , whenever she is not winning and $A > p$. If Bob follows eBay's guidelines, he will lose the auction; if instead he snipes--that is, bids his WTP only moments before the auction terminates---he can win.

Indeed, sniping is an equilibrium bidding strategy in eBay auctions [2]. Moreover, more experienced eBay bidders are more likely to snipe; experts bidding on eBay report that by sniping they can avoid revealing to other bidders private information about their WTPs. At the extreme, if all bidders were snipers, then eBay auctions would reduce to sealed-bid auctions, where all bidders submit their bids in private (for example, in sealed envelopes).

Amazon. Alternatively, Amazon implements the going, going, gone feature of live auctions. Whenever a bid is cast in the last ten minutes of an auction, the auction is automatically extended for an additional ten minutes from the time of the latest bid. This ensures that an auction can't close until ten bidless minutes have passed.

By design, sniping is not effective in Amazon auctions. Moreover, Amazon recognizes the value of creating a bidding frenzy: Amazon offers its sellers the option of enticing early bidding with a first-bidder discount. The intent is to start the bidding early, in an attempt to create a virtual atmosphere that is reminiscent of an English auction house, with lots of bidder excitement ultimately leading to greater revenues than can be achieved via sealed-bid auctions.

Which electronic auction design is superior? From the point of view of economic efficiency, where the goal is to maximize social welfare, the item should be awarded to the bidder with the highest WTP. eBay's auction design, then, is problematic because sniping can lead to inefficient outcomes (recall Alice and Bob). In contrast, Amazon-style auctions have been observed to be more efficient in laboratory experiments [3].

3. Sponsored Search Auctions

In sponsored search markets, search engines sell keywords, such as "digital camera," to advertisers. The goods in these markets are electronic billboards--spaces to post promotions alongside the results of keyword searches. As sellers, search engines seek to maximize their revenue; hence, they face the age-old question of how to price their goods.

Before the advent of the Internet, the most widespread advertising pricing model was based on the cost per thousand impressions (CPM), where an impression is a posting of an ad. This classic model, well suited to more traditional media such as television, newspapers, and magazines, was adopted in 1994 by Internet content providers. Historically, Internet CPM contracts were large (some early contracts were valued at \$30,000). Hence, their negotiation was cumbersome, necessitating extensive human intervention. This was unfortunate because, on the Internet, there is ample opportunity for automation.

Overture (Yahoo!). In 1997, Overture (now Yahoo! Search Marketing) launched an innovative framework for selling advertising space on the Internet. There were (at least) two novel aspects of Overture's design:

- Rather than selling large, expensive chunks of advertising space, each keyword was sold via its own auction (with prices as low as $< \$1$ for a single slot).
- Payment is made on a pay-per-click (PPC), rather than a per-impression, basis (where the advertiser pays only if someone clicks on the ad).

Suddenly, with Overture on the scene, combinatorially many inexpensive, perishable goods--namely, keywords in searches--were up for auction. For each keyword, several advertising slots were auctioned at once, each one representing a position relative to the top of the search page. Because of presentation bias (the tendency of users to click on earlier search results), positions closer to the top are preferred.

Overture's auction mechanism has been characterized as a generalized first-price auction (GFP) [4]. In a first-price auction for a single item, the

workshop that is focused on the design and analysis of single and multiagent systems, and how they can be modeled to support e-commerce

Workshop on Game Theoretic and Decision Theoretic Agents (GTDT): an annual workshop that focuses on developments in game and decision theory, as applied to agent-based systems

Reviews

Derandomization of auctions
Aggarwal G., Fiat A., Goldberg A., Hartline J., Immorlica N., Sudan M. *STOC 2005*, Baltimore, MD, May 2005, pp. 619-625

Combinatorial auctions
Cramton P., Shoham Y., Steinberg R., 2006

An agent-based approach to flexible commerce in intermediary-centric electronic markets

Liu D., Hwang T. *Journal of Network and Computer Applications* 27(1): 33-48, 2004

Exploitation vs. exploration: choosing a supplier in an environment of incomplete information

Azoulay-Schwartz R., Kraus S., Wilkenfeld J. *Decision Support Systems* 38(1): 1-18, 2004

Matchmaking in electronic markets : an agent-based approach towards matchmaking in electronic negotiations
(LNCS 2882) Veit D., 2004

Consumer search in electronic markets: an experimental analysis of travel services

Öörni A. *European Journal of Information Systems* 12(1): 30-40, 2003

TAC-03: a supply-chain trading competition

Sadeh N., Arunachalam R., Eriksson J., Finne N., Janson S. *AI Magazine* 24(1): 92-94, 2003

Decision support for managing an electronic supply chain

highest bidder wins the item at the highest price. In a GFP, multiple items are up for auction; the highest bidder wins the first item at the highest price, the second-highest bidder wins the second item at the second-highest price, and so on.

Another notable aspect of Overture's auction design was that winning bids were posted. This feature, however, proved problematic, resulting in instability, as we can see in the following example.

A typical Overture auction for three advertising slots, with bidding agents Alice, Bob, and Carol, might proceed as follows: Alice bids her WTP (A), while Bob and Carol bid their WTPs (B and C), with $A > B > C$. Alice wins slot one at the price of A , while Bob wins slot two at the price of B , and Carol wins slot three at the price of C .

Since winning prices are posted, Bob quickly realizes that a bid of $C + \epsilon$ would be sufficient to win slot two, and he decreases his bid accordingly. Immediately thereafter, Alice notices that a bid of $C + 2\epsilon$ would be sufficient to win slot one. She, too, decreases her bid accordingly.

But the game does not end there. Bob would prefer slot one to slot two, so he bids $C + 3\epsilon$ to outbid Alice. Alice follows suit, bidding $C + 4\epsilon$ to outbid Bob. This bidding war continues until Bob's WTP is reached, at which point Bob drops his bid back down to $C + \epsilon$ and the cycle begins anew.

Google. In 2002, Google launched its AdWords Select program. Google's auction mechanism has been characterized as a kind of generalized second-price auction (GSP) [4]. In a second-price auction for a single item, the highest bidder wins the item at the second-highest price. In a GSP, multiple items are up for auction: the highest bidder wins the first item at the second-highest price, the second-highest bidder wins the second item at the third-highest price, and so on. The GSP eliminates some of the instability of the GFP, because there is no need for the winner of the k th item to adjust his bid down to the $k+1$ st price; by design, he already pays the $k+1$ st price.

Google's auction mechanism introduced another new feature, which exploits the fact that advertisers bid (and pay) on a PPC basis rather than a CPM basis. Instead of allocating advertising slots in the decreasing order of bids, slots are allocated in the decreasing order of expected revenue. This revenue is computed as the product of the advertiser's bid and the advertiser's expected click-through rate—an estimate of how likely the advertiser's ad is to be clicked on.

In sponsored search auctions, there are numerous design dimensions, including: CPM versus PPC versus pay-per-conversion; GFP versus GSP; and rank-by-bid versus rank-by-revenue. Among the various combinations, Google's auction mechanism is one that explicitly seeks to maximize Google's revenue. But Google's auction mechanism lacks one basic property that has been identified by game theorists as desirable of auctions: incentive compatibility (IC). In an incentive-compatible auction, bidders are incentivized to bid their true valuations, no more (of course) and no less. It is well-known [4] how to implement an incentive-compatible generalized auction. The design of an incentive-compatible generalized auction that maximizes revenue is an open problem in game theory whose solution could have a great impact in the world of e-commerce.

4. Trading Agent Competition

Thus far, several popular electronic auction designs have been discussed from a game-theoretic point of view. I will now describe the Trading Agent Competition (TAC), a game-based platform for joint experimentation with intelligent software agent strategies and electronic market designs. TAC is an annual competition hosted by the Association for Trading Agent Research. The purpose of TAC is twofold: to evaluate agent strategies in a simulated electronic market, and to evaluate market designs, given a group of trading agents.

TAC agents are programmed traders. They operate autonomously in the

Keskinocak P., Goodwin R., Wu F., Akkiraju R., Murthy S. *Electronic Commerce Research* 1(1-2): 15-31, 2001

Economics of dynamic pricing in a reputation brokered agent mediated marketplace
Zacharia G., Evgeniou T., Moukas A., Boufounos P., Maes P. *Electronic Commerce Research* 1(1-2): 85-100, 2001

Establishing online trust through a community responsibility system
Ba S. *Decision Support Systems* 31(3): 323-336, 2001

The dynamics of mass online marketplaces: a case study of an online auction
Hahn J. *Human factors in computing systems*, Seattle, WA, Mar./Apr. 2001, pp. 317-324

Bargaining and market behavior
Smith V., 2000

Differential games in economics and management science
Dockner E., Jørgensen S., Van Long N., Sorger G. (Eds), 2000

Do electronic marketplaces lower the price of goods?
Lee H. *Communications of the ACM* 41(1): 73-80, 1998

Market-based control
Clearwater S. (Ed.), 1996

Computational organization theory
Carley K., Prietula M. (Eds.), 1994

marketplace, sending bids, requesting quotes, accepting offers, and generally negotiating deals according to market rules; humans do not intervene while negotiations are in progress. These agents face a challenging task. To play the market effectively, they must make decisions in real time in uncertain, dynamic environments. Successful agents rapidly assimilate information from multiple sources, forecast future events, optimize the allocation of their resources, anticipate strategic interactions, and learn from their experiences.

Presently, there are two running versions of the Trading Agent Competition: TAC Travel and TAC Supply Chain Management (SCM). In TAC Travel, travel goods--flights, hotels, and entertainment tickets--are sold in three types of electronic markets. In TAC 2000, hotels were sold in ascending auctions with a fixed clearing time, much like eBay auctions. As you can probably anticipate, most TAC 2000 agents behaved like snipers. (The designers of the TAC 2000 hotel auctions anticipated this outcome. Their attempted fix, however, failed.)

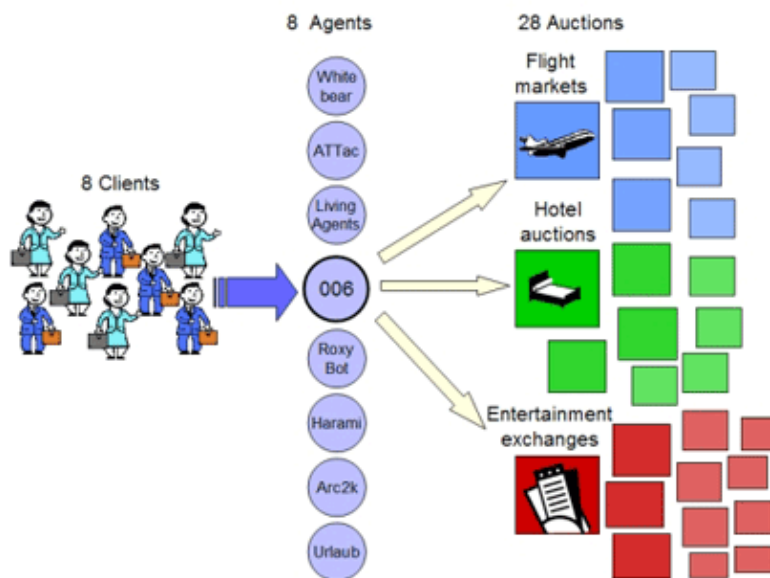


Figure 1: Illustration of the environment a TAC agent operates within. On the left are its eight clients and their preferences, in the middle are all of its competitors (seven competitors per game), and on the right are all of the auctions (28 simultaneous auctions of three different types).

(Illustration courtesy of the Swedish Institute of Computer Science (SICS) Intelligent Systems Lab.)

In TAC SCM, agents are computer manufacturers who buy raw materials from suppliers and sell finished goods to customers. Like TAC Travel, early renditions of TAC SCM were effective in revealing problems with certain electronic market designs. In TAC SCM 2003, the suppliers' pricing model was susceptible to manipulation by the agents. This model was redesigned twice. Finally, it appears to be no longer susceptible to (egregious) manipulation.

Today, TAC Travel and TAC SCM are both running strong. Next year, the TAC Board of Trustees plans to release a new game on electronic market design.

5. Conclusion

Electronic commerce promises to reshape the world's economy. There has been a recent surge of activity in algorithmic game theory, a new research area that aims to develop computationally efficient market mechanisms. Currently, however, many popular electronic marketplaces have their shortcomings. This is true of eBay's consumer auctions, Yahoo!'s and Google's sponsored search auctions, and the initial designs of TAC Travel and TAC SCM.

The design of electronic market mechanisms is a tricky business. Electronic

market mechanisms can (and should) be evaluated based on game theory, prior to their release. Populating simulated markets with artificially intelligent software agents in a forum such as the international Trading Agent Competition is another indispensable means of evaluation. This research area, which overlaps with both game theory and computer science, is sometimes called computational mechanism design.

The interplay of game theory and e-commerce is an exciting domain for future research. Progress in this area will require a combination of theoretical analysis, empirical studies, and simulation experiments. Better market designs will do a better job of matching buyers with sellers, ultimately enhancing the welfare of our society.



Created: Jul 28 2006
Last updated: Aug 7 2006

References

- 1) Roth, A.E., Ockenfels, A. Last-minute bidding and the rules for ending second-price auctions: evidence from eBay and Amazon auctions on the Internet. *American Economic Review* 92, 4 (2003), 1093–1103.
- 2) Ockenfels, A., Roth, A.E. Late and multiple bidding in second-price internet auctions: Theory and evidence concerning different rules for ending an auction. *Games and Economic Behavior* 55, 2 (2006), 297–320.
- 3) Ariely, D., Ockenfels, A., Roth, A.E. An experimental analysis of ending rules in Internet auctions. *Rand Journal of Economics* 36, 4 (2005), 891–908.
- 4) Edelman, B., Ostrovsky, M., Schwarz, M. Internet advertising and the generalized second price auction: selling billions of dollars worth of keywords. NBER Working Paper No. W11765, 2005.

REVIEWER'S AREA

MASTHEAD

SUBSCRIBE

PRESS

TIPS

HELP

CONTACT US

Reproduction in whole or in part without permission is prohibited. Copyright © 2000-2008 Reviews.com
[Terms of Use](#) | [Privacy Policy](#)