Economic Agents for Automated Trading

Chris Preist
Living the Vision
HP Laboratories Bristol
HPL-98-77
April, 1998

E-mail: cwp@hplb.hpl.hp.com

agents, ebusiness,
internet, trading

Economic Agents for Automated Trading

Chris Preist
Agent Technology Group
Hewlett Packard Labs
Filton Road
Bristol BS12 6QZ

cwp@hplb.hpl.hp.com

1. Commerce meets the Web

With the advent of the world wide web, it is becoming possible to trade with organisations and individuals across the world at the click of a mouse. This trade doesn't just take place between consumers and businesses - the bulk of it is between businesses. General Electric alone engages in internet transactions with its suppliers worth over $1Bn a year – almost twice the total amount of consumer sales on the web [ECO 97]. By the year 2000, it aims to buy all its industrial supplies electronically [ECO 98]. Industry analysts expect more and more trade to take place in this way, as businesses wake up to the speed, flexibility and ease of global access the net offers.

Negotiation plays an essential part of most business-to-business transactions. Currently, representatives of the companies involved must do this by phone or email, even when the trade is taking place on the internet. The time and effort to do this means they only negotiate with a small number of possible trading partners, missing good deals elsewhere. However, we expect the world to be very different in the near future. Responsibility for much of this negotiation will be handed over to software agents. These agents will monitor other trade agents continuously, watching for potential opportunities. They will be able to enter into negotiation with many potential trade partners at once, reaching an acceptable deal and setting up a contract in a matter of milliseconds.

As the time and cost of making a contract drops rapidly, the nature of contracts will change. Contracts between businesses will no longer be things laboriously set up, lasting for months or years. Instead, they could last for as little as a single transaction. New electronic marketplaces, trading goods such as bandwidth on transatlantic cable connections, will come into being. Agent technology will form a central pillar of this new world of business.

2. Agents in Electronic Commerce

Agents are already playing an important role in electronic commerce. Guttman et. al. [Gutt 98] have shown how agent technology can contribute to different aspects of consumer buying – deciding what to buy, who to buy it from, how much to pay, and finally the actual exchange of money and goods.

Agents are now actively used in e-commerce to search the web and find the trader selling a given product at the cheapest price. BargainFinder [BF] was the first of these, and received a mixed response from sellers. Some web-based trading sites initially wanted to ban it, out of fear that it would drive prices down, as customers would use price alone to determine where to buy from. Other sites welcomed it in, as it would bring them new customers. Jango [JG] is also an agent which searches for the best deal. It appears as a web-browser to the site being visited, so sellers don't know if it is an agent or a person checking out their prices. As a result, it is harder to block.

Other agents help a customer determine exactly what it is they wish to buy. Firefly [FF] compares a user's taste in music or film with a large database of other peoples preferences. It recommends that the user try products which are highly rated by other people who have similar preferences. PersonaLogic [PL] helps a user select the best model of a product, such as a car, for their needs based on a series of questions and answers.

Internet auctions are becoming an increasingly popular way of determining the price at which a good is sold at. There is a simple agent, SmartBidder [SB], that will bid on your behalf in such an auction. AuctionBot [AB] allows users to write their own negotiating agents, or to bid themselves.

However, all these are focussed on business-to-consumer transactions. There are currently no agents on the Internet which focus on business-to-business trading. This is because such transactions require negotiation, and solving the problem of automated negotiation is not easy.[1]

3. Agents which Negotiate

If an agent is to negotiate on behalf of an organisation, it needs:

- A representation of the goods or services which are to be traded.
- An understanding of what the organisation wishes to achieve from the negotiation.
- A strategy for negotiation which is at least as effective as a qualified person in the same situation.

Representing the goods to be traded can be complex. It is important to develop a formal representation which makes it possible to specify what a product is, and also how it differs from similar products made by other manufacturers. What features does it have? How well does it perform? This must be done in a way that is perceived as fair by all businesses involved.

The goals of an organisation during negotiation can often be very subtle and ill-defined. A buyer may be willing to pay more for a higher quality product, or for an added feature. Such tradeoffs are often made instinctively. Hence it is very difficult to capture the exact criteria behind such decisions, allowing an agent to automate them.

However there is a class of goods, *commodities*, for which these two factors are less of an issue. A good is a commodity if price is the only factor that is considered when trading it – it cannot be differentiated by being of superior quality, or by adding extra features. Goods such as crude oil, electricity and wheat are all commodities. Other goods, such as memory chips, network bandwidth and even personal computers, are close to being commodities. Because price is the only factor involved in comparing one potential deal with another, it becomes easier to represent the goods traded and the organisation's goals. For this reason, the first deals to be negotiated by automated agents will be for commodity-like goods.

4. Negotiating Strategies

Given that agents know what goods they are to negotiate about, this brings us to the third factor listed above: what should the negotiating strategy be? This is a hot topic of research in the agent community. There are three main approaches currently being explored:

- The rule based approach: A negotiation strategy is coded into the agent using a set of rules specifying how they should bargain, and when they should or shouldn't accept a deal. (For example, [Sierra et. al. 97])
- The game theoretic approach: Economic analysis of a negotiation problem is used to propose an appropriate negotiation protocol and strategy. ([RZ 94],[VJ 98])
- The adaptive behaviour approach: Simple agents adapt their strategy by observing the behaviour of the marketplace and their current performance. ([Tesf 97],[CB 98])

Research in the first two areas has so far focused on negotiation between two parties only, or on appropriate strategies in an auction with one seller. We believe that much of the negotiation that will take place on the internet will involve many possible buyers and sellers, all negotiating with each other simultaneously. We at Hewlett Packard would like to create trading communities which are highly efficient, well-behaved, and in which each participating agent negotiates at least as well as a human would in the same situation. To do this, we have chosen to take the adaptive behaviour approach to agent design. The work presented here is an extension to that presented in [CB 98].

5. Marketplaces where Agents Meet

As we are interested in markets with many buyers and sellers, simultaneously negotiating with each other, we have focussed our work on the *double auction* (DA) market [Fried 92]. In a DA, buyers announce *bids* to buy goods at a

[1] Consumer-to-consumer trading can also involve negotiation. In this paper, we will focus primarily on business to business negotiation. See [CDGM 97] for an example of agent-based negotiation between consumers.

given price. Similarly, sellers announce *offers* to sell goods at a given price. These bids and offers can be made at any time, and are heard by all participants. Participants are free to accept any bid or offer that they like, or to announce new bids and offers. In this way, multi-party negotiation takes place. Traders adapt their pricing in response to what they observe others doing. The double auction evolved from informal gatherings of sellers (such as wheat farmers) with buyers in local markets, and is now a well-established mechanism used in international commodities markets.

When traders participate in a double auction market, an interesting emergent phenomenon occurs; the price at which trades take place tends to stabilise at a certain value. This value is known as the *equilibrium price*, and is determined by the law of supply and demand; the equilibrium price is the price at which the number of goods offered for sale is equal to the number of goods the buyers wish to purchase at this price. At equilibrium, the maximum amount of trading takes place, and the traders benefit accordingly. If the supply of the good and the demand for the good remains constant, then trades will continue to take place at this price. However, if the supply or demand alters (maybe because a seller reduces production costs, or because more buyers appear), then the equilibrium price also alters. Trading following such a change will initially take place at various prices, but will soon settle at the new equilibrium. (For more on supply and demand, and how it determines market price, see for example [Silv 95])

An agent-based marketplace should also have this property. The agents should converge on the equilibrium price at least as quickly as humans do in the same situation, and should respond to changes in supply and demand by rapidly moving to the new equilibrium price. We have successfully used simple adaptive agents to develop an experimental automated marketplace that does exactly this.

We have chosen to focus on a particular style of double auction marketplace, the *persistent shout double auction*. In this setup, a trader may make a bid or offer at any time, but once made it persists until the trader chooses to alter it or remove it, or it is accepted. One example of such a marketplace exists on the internet: Fastparts [FP] provides a persistent shout double auction for buying and selling excess electronic components. Buyers and sellers place bids and offers on a web-based trading floor. They revise their bids/offers in response to other trading activity. When a bid and offer meet at the same price, they are deleted and a trade takes place at that price. The New York Stock Exchange also uses a form of persistent shout double auction; the NYSE rule states that the current bid and offer persist, and that any new bid or offer must improve on the existing one. However, unlike the Fastparts marketplace, a ‘reset’ occurs when a trade is made, and previous shouts must be repeated.

6. An Agent Based Double Auction

In our current marketplace, agents buy and sell an abstract commodity from each other. In the future, this commodity could represent grain, memory chips or communications bandwidth. Agents are divided into buyers and sellers, with each agent wishing to trade one good in a given trading ‘day’. Each agent is given its own limit price; if it is a buyer, it will never buy for over this price, and if it is a seller, it will never sell for less than this. They are free to make any bid/offer subject to this constraint, and prefer to make a trade at their limit price than to not trade at all. Time is divided into trade ‘rounds’. In the first round of a day, all agents must shout their initial bid or offer. In subsequent rounds, any agent can modify their current bid/offer if they choose. If they choose not to, then their current bid/offer will remain. As with the Fastparts marketplace, we assume a trade takes place if a bid and offer meet at the same value. If a bid and offer cross, (i.e. a bid is made which is higher than an offer,) we assume the trade takes place at the average of the two prices. If more than one bid and offer cross, then multiple trades can take place in the same round: The highest bid is matched with the lowest offer, these are deleted and the process is repeated until there is no overlap between bids and offers. Agents continue trading until all agents have bought/sold or are no longer willing to adjust their bid/offer. At this time, the trading ‘day’ is over, and all agents are reinitialised with an intention to buy or sell one good.

7. Supply and Demand

The limit prices given to the agents determine the underlying supply and demand curves. For example, consider an experiment with 5 buyer agents and 5 seller agents. Let the buyer agents $b_1,\ldots,b_5$ be given limit prices of $0.50, $1.00, $1.50, $2.00 and $2.50 respectively. Similarly, let the seller agents $s_1,\ldots,s_5$ also be given limit prices of $0.50, $1.00, $1.50, $2.00 and $2.50 respectively. This means that if the good is being traded at $1.00, then buyers b_2, b_3, b_4 and b_5 each wish to buy one unit in a day, and hence the quantity demanded is 4 units. Similarly, sellers s_1 and s_2 wish to sell, and hence the quantity supplied is 2 units. In this way, we can calculate the quantity supplied and demanded at different prices. We can plot this information on a graph, to give two curves; the *supply* curve and the *demand* curve. (Figure 1).

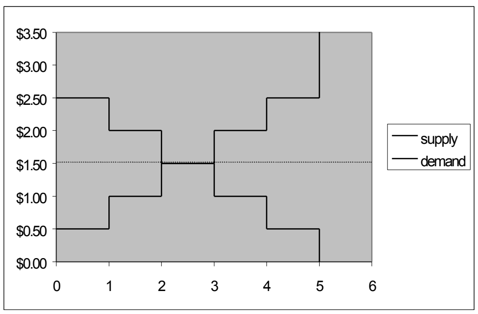

Figure 1: Supply and Demand curves

Where these two curves intersect, the quantity supplied is the same as the quantity demanded, giving the equilibrium price. In this case, the curves intersect at $1.50. At this price, either two or three goods will be traded, depending on whether traders s_3 and b_3 choose to trade and make no profit. In our experiments, traders will trade for zero profit in preference to not trading at all.

8. Agents which Adapt to the Market

The algorithm the agents use consists of a small number of common-sense heuristics combined with a simple learning rule. Each agent keeps track of the profit it currently hopes to make. A buyer agent will be willing to pay up to its limit price less this profit goal. A seller will be willing to sell for any price greater than its limit price plus its profit goal. The price an agent is willing to trade for is its *current valuation*.

Initially, each agent is assigned a positive random profit goal. Each agent then monitors bids, offers and trades in the marketplace, and modifies its profit goal so as to maximise profit. If it sets its profit goal low, it will not make as much profit as if it sets its profit goal high. However, if it sets its profit goal too high relative to the market, it will fail to make a trade. The agent must use information about current market activity to find the balance, and must respond to changes in the marketplace if a new balance is appropriate.

The algorithm runs each market round and consists of two phases. Firstly, a small set of heuristics uses current market activity to determine a target trade price. Then a simple learning rule is used to determine how much the current valuation (and hence the profit goal) is altered towards the target.

The heuristics are:

```
For BUYERS;

If the highest bid is below the
lowest offer, then target just
above the highest bid.

If the highest bid is equal or
greater than the lowest offer,
then target the lowest offer.

For SELLERS;

If the highest bid is below the
lowest offer, then target just
below the lowest offer.

If the highest bid is equal or
greater than the lowest offer,
then target the highest bid.
```

These rules may result in an agent either reducing or increasing its profit goal. If an agent has a good to trade, it should be willing to reduce its profit goal if it is failing to make a sale because of its competitors. However, if an agent currently has no good to trade, it should not reduce its profit goal, as it has no reason to enter into competition with agents currently trading. For that reason, we place the additional constraint that if the above rules require such an agent to reduce its profit goal, then it does not adjust its valuation this round.

The intuition behind these heuristics is straightforward. If trades are not taking place, an agent should attempt to be the most competitive by making the best bid/offer, so should target a valuation slightly better than its competition. If, on the other hand, trades are taking place, an agent should target the best price at which it can obtain a trade.

The heuristics determine what value, if any, the agent should target in a given round. The learning rule then determines how far the agent adjusts its current valuation towards this value. The learning rule used is *Widrow-Hoff*, which is also used for back propagation learning in neural networks [eg RHW 86]. It is parameterised by a *learning rate* b which determines the speed with which the adjustment takes place. b is a value between 0 and 1. The higher b is, the nearer the current valuation moves to the target.

9. Stable Marketplaces of Agents

Given this experimental setup, the agents rapidly adjust their profit goal to take account of current market conditions. Without any explicit representation of supply, demand or equilibrium price, they trade at equilibrium price.

Figure 2 shows the adjustment taking place. This shows an experiment run with 11 buyers and 11 sellers, where the predicted equilibrium price is \$2. We plot the actual price of trade against the time the trade took place. We can see that in the first day, trades take place at prices between \$1.73 and \$2.16. On the second day, the variation of price is reduced – the price ranges from \$2 to \$2.15. On subsequent days, this trend continues and trades take place increasingly close to the equilibrium price.

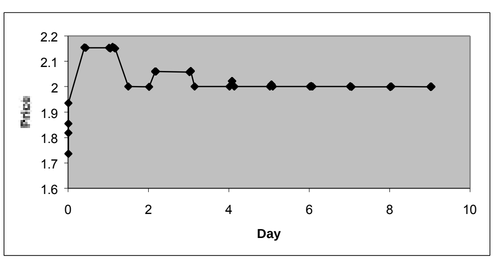

Figure 2: Time series plot of trades

To study the convergence to equilibrium, we use *alpha*, a measure introduced in similar experiments with human traders by experimental economist Vernon Smith [Smith 62]. Alpha is defined to be the standard deviation of the actual trades in a given day around the equilibrium trade price, expressed as a percentage of this price. Hence, if alpha is small, it means trades are taking place close to equilibrium. We run the above experiment 50 times, and plot the mean value of alpha each day in figure 3. The dotted bars either side give the mean value plus and minus one standard deviation. We see that on the first day, alpha is around 5%, but on the second and subsequent days, it is under 1%.

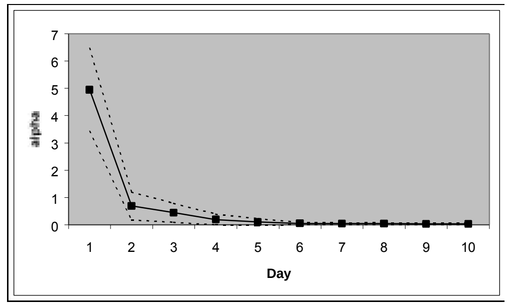

Figure 3: Convergence to equilibrium

This performance is significantly better than humans. For an identical experimental setup, human traders gave an initial alpha value of 11.8%, reducing to 3.5% after 5 days trading [Smith 62].

If the community of agents increases in size, convergence to equilibrium takes place more rapidly. For example, in an experiment with 501 buyer agents and 501 seller agents, alpha is under 4% after the first day, and is at 0.1% after 3 days.

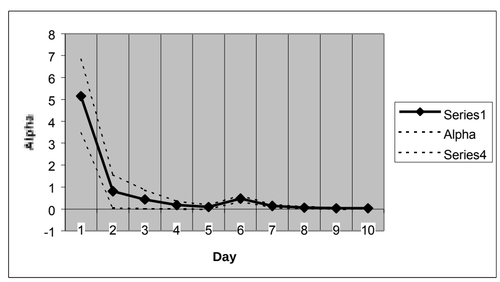

Figure 4: Handling a change of supply

The agents are able to handle a sudden shift in supply or demand, with only a small disturbance in alpha. After a brief period of non-optimal trading, they move to the new equilibrium price. Figure 4 shows the plot of an experiment in which supply decreased in day 6. Again, we use 11 buyer agents and 11 seller agents. The initial equilibrium price is $2, but after the decrease in supply, it becomes $2.75. Alpha increases from 0.1% to 0.6% on day 6, but soon reestablishes stability at the new equilibrium price.

Our agents are efficient, lightweight programs that rapidly move to trade at equilibrium price. They produce more stable trading behaviour than humans under a similar experimental setup. They therefore reach deals that human traders would lose, even though the trades are in their interests.[2] For this reason, we believe that they can form the basis of effective electronic marketplaces on the internet.

10. Moving to Realistic Markets

Inevitably, experimental investigations make simplifying assumptions. The market we have used to carry out our work is idealised. If the work is to be used in anger, these idealisations and simplifications must be overcome:

- The current system is based around the concept of a trading 'day'. All agents wish to trade exactly one good each 'day'. In reality, the intention to buy or sell goods can arrive at any time, asynchronously. Agents must be able to handle this.

- Agents currently trade only one good at a time. In real marketplaces, sellers may be selling several of the good, and buyers may need to purchase more than one.

- We cannot assume that supply and demand are stable for a period, and then jump to new stable values. Supply and demand can be constantly fluctuating as circumstances change. Agents must be able to track these changes, and exploit them.

- Agents may be under time pressure to close a deal [VJ 98] (For example, because they are purchasing goods to be used in manufacture.) They must adjust their pricing strategy to take this into account.

- The behaviour of an agent is determined by its limit price. It is not always easy for an organisation to determine exactly what its limit price is. An agent should instead be able to make use of whatever information is provided by the organisation, and convert this into appropriate economic representations, such as limit price.

We are currently addressing these four issues in our research in agent-based commerce. By solving these, we hope to make automated trading on the internet an everyday reality. Specialised electronic marketplaces will emerge, where buyers and sellers can trade remotely and automatically. Some of these markets will trade raw materials, such as paper and metal. Others will trade overstocked or otherwise unwanted goods, such as electronic components. Others will trade virtual goods, delivered electronically, such as communications bandwidth and access to information feeds. Others will trade commodity-like services, such as language translation or remote contract programming. An agent will constantly monitor these marketplaces for goods or services of interest to its company. If it sees something of interest, it will enter into high speed negotiation with agents representing the sellers, and will make the purchase if a mutually acceptable price can be reached.

Initially, these marketplaces will only handle commodity-like goods. However, as agents become more sophisticated in their representation of goods and needs, and their ability to negotiate, more complex products could be traded automatically – computer systems, catering contracts and hotel booking, to name a few.

Agents will not only negotiate. They will also play an active role in all aspects of electronic commerce between businesses, from the decision of what to purchase to after-sales support. They may determine when to initiate purchases, by automatically monitoring inventories and making predictions about future stock requirements. They may keep track of customer accounts, determining which new products customers are likely to be interested in. They may monitor a product after sale, and inform the supplier when it needs servicing.

Agents are destined to play an increasingly central role in electronic commerce in the 21st century, and research into automated negotiation is central to making this happen.

[2] This is because trade is taking place away from equilibrium price, so fewer trades take place. Someone who should gain from trading therefore loses out.

References

[ECO 97] *Electronic Commerce Survey*. The Economist, May 10, 1997.

[ECO 98] *To byte the hand that feeds it*. The Economist, Jan 17th, 1998.

[Gutt 98] R. Guttman, A. Moukas and P.Maes. *Agent Mediated Electronic Commerce: A Survey*. Knowledge Engineering Review, June 1998.

[JG] Jango URL:
http://www.jango.com/

[FF] Firefly URL:
http://www.firefly.com/

[PL] PersonalLogic URL:
http://www.personalLogic.com/

[SB] SmartBidder at Numismatists online URL:
http://www.numismatists.com/

[AB] AuctionBot URL:
http://kasbah.media.mit.edu/

[CDGM 97] A. Chavez, D.Dreilinger, R.Guttman and P. Maes. *A Real Life Experiment in Creating an Agent Marketplace*. Proceedings of the Second International Conference on the Practical Applications of Agents and Multi Agent Systems, 1997.

[Sierra et. al. 97] C. Sierra, P. Faratin and N. Jennings. *A service oriented negotiation model between autonomous agents*. Proc. Modelling Autonomous Agents in a Multi Agent World, 1997.

[RZ 94] J. Rosenschein and G. Zlotkin. Rules of Encounter. MIT Press.

[VJ 98] N. Vulkan and N. Jennings. *Efficient Mechanisms for the Supply of Services in Multi-Agent Environments*. Technical report, Dept. of Electronic Engineering, Queen Mary & Westfield College, University of London

[Tesf 97] *How Economists can get Alife*. In The Economy as an Evolving Complex System II, Addison Wesley 1997.

[CB 98] D. Cliff and J. Bruten. *Less than Human: Simple adaptive trading agents for CDA markets*. To be published in Proceedings of CEFEES ’98.

[Fried 92] *The double auction market institution*. In The Double Auction Market: Institutions, Theories and Evidence. Friedman and Rust (eds) Addison-Wesley 1992.

[Silv 95] E. Silverberg. Principles of Microeconomics. Prentice-Hall 1995.

[FP] Fastparts URL:
http://www.fastparts.com/

[Smith 62] V. Smith. *An Experimental Study of Competitive Market Behaviour*. Journal of Political Economics, 70, 111-137.

[RHW 86] D. Rumelhart, G. Hinton and R. Williams. *Learning internal representations by error propagation*. In D. Rumelhart and J. McClelland. Parallel Distributed Processing, Volume 1: Foundations. MIT Press 1986.