

Economic Implications

of

Agent Technology and E-Commerce

*Nir Vulkan*¹

July, 1998.

Abstract

Following the exponential growth of the Internet and World Wide Web, electronic commerce already accounts for some \$1.5 billion a year. But the scope and volume of e-commerce is likely to increase further sharply because of advances in agent technology, i.e. applications where users delegate the authority to search and filter information, schedule meeting or negotiate agreements, to *software agents*, programs that acts independently on their behalf.

This paper describes in some detail existing agents and multi-agents applications in the context of e-commerce, and suggests a research agenda for economists in response to these changes in technology and lifestyle. First, several ways where economic theory (and in particular implementation theory) can be used to design and improve the efficiency of e-commerce systems are described. Second, the paper discusses the impact on markets of using software agents. Finally, the paper discusses how economic theory can be used towards the design of the interactions between agents and their users.

1. Introduction

¹ Department of Economics, University of Bristol, 8 Woodland Road, Bristol BS8-1TN, U.K., and The Centre for Economic Learning and Social Evolution (ELSE), University College London. E-mail: N.Vulkan@bristol.ac.uk

This paper is, in parts, based on my lectures at HP labs (Bristol), The Computer Science department at Tel-Aviv University, and at the Technion, the Israeli Institute of Technology (Haifa). I am grateful to Ian Jewitt, Pasquale Scaramozzino, Christian Schluter, and participants of the HP workshop on Internet economics for helpful comments and discussions.

The exponential growth of the Internet is the single most important development in Information Technology in the last decade. With Internet traffic doubling in less than every four months, and WWW traffic increasing more than 15 times in 1995 alone (Quarterman, 1995), there are no signs that this phenomenal growth rate is slowing down. The Internet offers 24 hours a day access to a large number of potential customers from all over the world with relatively low overheads. It is therefore not surprising that Internet-based commerce is already emerging as a significant phenomena, with some companies, like Amazon and Microsoft's Expedia, which sell only via the Web. Most of the Internet-based commerce falls under the following two categories: First, on-line shopping, where customers search the WWW using a Web browser (like Netscape or Internet Explorer) and a search engine (like Yahoo!). These type of interactions can be characterised as *user-driven* transactions, and are sometimes refer to as *first-generation e-commerce*. The second, far less publicised, component of e-commerce consists of the business-to-business interactions which are carried out by large organisations (for example trading with suppliers). In the last year alone, General Electric's turnover of Internet based trade was in excess of \$1 billion, twice the total of world-wide Internet sales in the same year.

Despite these impressive figures, the overall feeling amongst practitioners is that we have yet to experience the full impact of electronic-based commerce. But this may change soon: *Agent technology*, which is already affecting almost every aspect of computing, could become for e-commerce what Windows was for PCs - a relatively simple and user-friendly way of utilising the new technology. In this framework individuals and organisation interact via the network using *software agents*. A software agent (also *automated*, or *autonomous* agent) is a program that acts independently on behalf of its user, in furtherance of its interests. Moreover, some of these agents are capable of copying themselves over the Internet, of interrogating host Web sites, and of interacting with other agents. Unlike first-generation e-commerce systems - which are likely to account for the bulk of electronic commerce for the immediate future - the

The views presented in this paper do not necessarily represent the views of any of the commercial organisations mentioned in this article.

main characteristic of a second-generation systems is that the user *delegates* the authority to transact business to the agent.

Agent technology is hailed by many as the new revolution in software. The market research firm *Ovum* predicts a \$4 billion software agent market by the year 2,000 with wide-spread applications to telecommunications, marketing, entertainment and military market segments (Guilfoyle and Warner, 1994). The move towards agent-based environments is supported on the one hand by rapid progress in communication languages (like KQML and Telescript) and standardisation of electronic interactions, which means that agents can now securely identify themselves and carry out transactions, and on the other hand by an explosion of applications which are agent based (this is known in computer science as the agent-based programming approach, see Jennings & Wooldridge (1997)). Section 2 explains in more details the nature of these developments and lists some of their important applications.

It is easy to see why these agents have generated so much interest and hope for further growth of Internet based commerce: agents significantly increase the ability of users to search the Internet, and to sort out the outcomes of such searches according to users' preferences. Moreover, more advanced agents are now available which are also able to perform interactive searches on behalf of their users (for example, an agent searching for airline tickets from virtual travel agencies on the Web can match preferred dates, price-range, class of travel, and other features of the journey, without having to go back to its user at any given stage). Once these agents are equipped with some negotiating skills, they can be used to schedule meetings on behalf of their users, participate in on-line auctions, and trade in financial markets² (See the first chapter in Rosenchein & Zlotkin (1994), for a comprehensive discussion of possible future applications for negotiating software agents). By delegating the search and matching activities to automated agents, the comparative advantage of the Internet over traditional forms of business becomes significant. Once users begin to use these agents

² Computerised trading programs, currently banned from trading after several large market swings, could be re-introduced once the problem of software homogeneity ceases to exist, because trading agents will more closely resemble the characteristics of their users.

on a large scale, it is likely that the volume of Internet based commerce will sharply increase.

The purpose of this paper is to point out a research agenda for economists in response to these changes in technology and lifestyle, which is partially based on my experience of working with computer scientists over the last several years. With the notable exceptions of research on congestion and pricing of networks (see, for example, Mackie-Mason & Varian (1995), or Huberman & Lukose (1997)), it is disappointing how little work has been carried out by economists in response to these developments³. This is true not only because of the scale of these developments, but also because some of the issues raised are directly related to the types of problems traditionally studied by economists. In many cases, as shown in this paper, the tools of economic theory are best suited to address these issues. More specifically, three research categories are suggested, each dealing with a different aspect of e-commerce which could be addressed by economists. These categories are:

1. The design of markets for automated interactions, of protocols for multi-agent interactions, and of negotiating agents.
2. The impact on markets of using software agents, and
3. The interactions between users and agents.

Economic theory had, for many years, looked at the relationship between market structures and efficiency of outcomes. The design of an automated market for self-interested agents is that respect no different from any other market. Companies who set up electronic markets, on the Internet or elsewhere, are seeking to design these markets in such a way that will maximise their future profits through efficiency and competitiveness. The designers of multi-agent systems are choosing negotiation protocols⁴ on much similar criteria. In the long run, the success of these markets and systems will depend on the performance of self-interested agents, the behaviour of

³ Moreover, users' overwhelming objection to any changes in the way the Internet is run, including the idea of pay-as-you-use for the Internet, suggests that the research into pricing of Internet traffic is unlikely to be put to practical use.

which cannot be controlled. Accumulated knowledge from economic theory and mechanism design can therefore be put to use towards these goals. Over the last three decades, rapid progress in implementation theory (mechanism design), had brought the subject to an engineering-like state, where a large number of well understood mechanisms can be prescribed for participants with a given set of preferences. Trouble is, as the recent signalling accusation in the much publicised FCC auction show, human agents manage to always find ways to outsmart the designers of these mechanisms. Markets for automated agents prove therefore to be a much more suitable application for implementation theory. An automated agent is a pre-programmed algorithm, very much like the game-theoretical concept of strategy. In this sense, game theory (and mechanism design) seems much more suitable for automated agents than it is for humans (a point recognised by computer scientists, like in Rosenchein & Zlotkin (1994)). Section 3 shows how the intuitions gained from game theory can be put to use in the design of automated markets, protocols for multi-agent interactions, and the design of the agents themselves.

The importance of the second category suggested above should be relatively clear. This is essentially an economic problem: Search agents increase consumer's search power, and in general are thought of as increasing competitiveness in markets (at least markets for homogenous goods). However, early experiments (notably Anderson's experiment with BargainFinder, a CD search agent) show that other effects might also be present, and that the overall effect is unclear. Section 4 explains the need for a general modelling framework, based on the underlying incentive structure of all participants. A better understanding the incentive structure of electronic commerce has the additional advantage that it can be used as a basis for any future regulatory and/or taxation schemes.

The relationship of the third category to economics might not seem clear at first. The interaction between the user and its agent is clearly crucial for the success of agent based applications. A user must learn to trust her agent to act in her best interests if

⁴ The negotiation protocol for software agents is very similar to what economists know as the mechanism. An English Auction or an alternating offers protocols are two examples which are already being widely used for automated agents.

she is to empower it to make decisions, especially financial decisions, on her behalf. But there is more to this interface than what is normally subscribed by the science of software engineering and interface design. This is because the consumer must be able to express her utility function to her agent. This is particularly important when the object or service sought after is multidimensional (like hotel accommodations or travel arrangements). Of course people are not always aware of their preferences, let alone they are able of expressing them in a mathematically consistent form. The burden of constructing preferences and utility functions from past behaviour and by asking the right questions, will therefore fall on the shoulders of those who design agents. Section 5 describe how economic wisdom can be used towards future research into the design of software agents.

Some of the issues discussed in this paper are still theoretical: In many cases the technology is not quite there yet. Many of the applications discussed rely on unambiguous identification and on secure transactions, which are legally and technically not fully resolved yet. But since efforts are currently being made by governments and firms to overcome these hurdles, it is now worthwhile to understand the unique features of electronic marketplaces for self-interested electronic agents.

The rest of the paper is organised in the following way: section 2 provides some technical background for those readers who are not familiar with agent technology, and also lists some of the existing applications. Readers who are familiar with these developments can skip section 2. Section 3 describes current and future work into the design of multi-agent systems and artificial markets. Section 4 describes the challenge to modellers considering the effect on markets of using software agents. Section 5 deals with some of the economically relevant issues of user-agent interactions. Section 6 concludes.

2. Technical background

Agents are computer systems capable of acting autonomously, without the direct intervention of users. Like its human counterpart (for example a travel agent), a

computerised agent is entrusted to perform a task on behalf of someone. Most agents possess some degree of at least one of the following properties:

1. *Mobility*: A computer code is *mobile* if it is capable of making copies of itself, from one site to another, over a network. Computer viruses are early examples of “bad” mobile codes. Search agents, like BargainFinder and Jango, are examples of less harmful mobile codes.
2. *Intelligence*: the capability to interpret, learn and improve. Every form of intelligence which could improve the agent performance, excluding social intelligence (i.e. strategic considerations).
3. *Agency*: The ability to interact with other agents, ranging from “naive” to “strategic” exchange of messages. For example, negotiation skills.

The following diagram, adopted from an influential IBM white paper (Gilbert et al. 1995), illustrates the relationship between these three attributes:

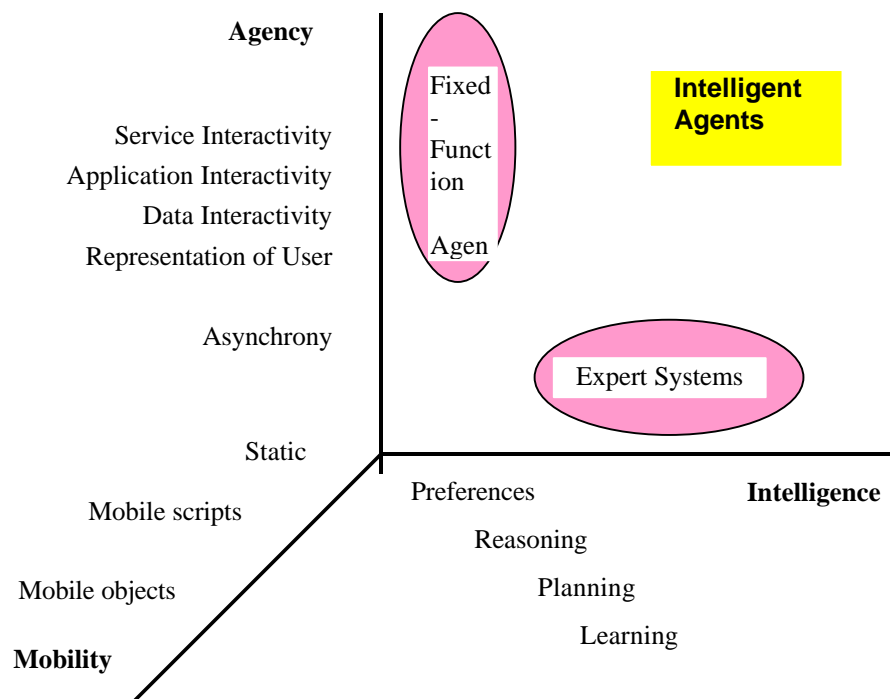


Figure 1: Scope of intelligent agents (Based on Gilbert et al. 1995)

Search agents are the more well known face of agent technology (not least because of the large sums of money recently paid by Microsoft and Excite for start-up companies based around such applications). Here the focus is on the performance of the single agent. Continuing with the terminology of the IBM white paper, search agents are mobile agents who are capable of interrogating host sites (mostly on-line shopping malls, for example Amazon.com) in search for a pre-specified pattern, like a price or a key word. The effectiveness of a search agent depends both on the size of the search, and the ability to get the relevant information from the host site (see Pazgal & Vulkan (1998) for a more detailed discussion of the performance attributes of search agents). The usage of this type of agent technology is widely expected to increase and follow a similar path to that of graphical user interfaces (GUI) in the 1970's; first it will be an optional extra, but over time products without intelligent agents will no longer be viable (Gilbert et al. 1995).

Multi-Agent Systems (MAS) are systems where two or more agents (human or software) interact via a commonly known protocol. Here, the focus is usually on the overall performance of the system (or the marketplace), rather than on the performance of a single agent. A software agent operating in a MAS is required, in addition to the properties described in the previous section, to:

1. use the same objective language as all other agents in the same MAS, and
2. possess some degree of negotiation skills.

More specifically, by an objective language we mean that there exists a set of commonly known symbols, rich enough for agents to be able to express their various goals and intentions (or, in the terminology of game theory, the dimensionality of the set of messages should be at least as large as that of the set of strategies), and where all agents' interpretation of these symbols is identical. These type of languages are therefore different from natural languages, like English or French (because those are ambiguous), but also different from regular programming languages, like C++ or FORTRAN (because those are not designed for expressing intentions and goals between different computer codes). Instead, the legal language used in courts may be a more suitable metaphor.

Agents operating in MAS are capable of communicating their intentions to one another, sharing limited resources and negotiating agreements. Currently, MAS are mostly used by large organisations (such as British Telecom, Boeing, and Mitsubishi) as internal resources managers. The reason why there are not any open-for-all MAS, or electronic marketplaces is the lack of sufficient standards and legislation. Efforts to overcome this hurdle are already taking place in academia and industry. Full scale, multi-agent systems where users can do business using negotiating agents, are therefore widely expected to become available over the next 2-5 years.

2.1. Applications

Agents are being used in almost every field of programming (so much that this is sometimes referred to as the “agent-based programming revolution”). Applications include e-mail filtering systems, digital library management systems, applications for patient monitoring, health care management, and entertainment. This year saw the third annual international conference on the practical application of intelligent agents and multi-agent systems (London, UK), with increasing numbers of participants from the industry. Because of limitation on space, only some of these applications are described below. The reader is referred to Jennings & Wooldridge (1997) for a (then) comprehensive survey of agent applications.

2.1.1 Internet Search Agents

The first and best-known example of a search agent is Andersen Consultants’ BargainFinder, which takes a request for a music CD from a user and searches a number of on-line CD shops for the best deal possible (see <http://bf.cstar.ac.com/bf/>). A second well known example is the Seattle-based agent Start-up Company NetBot, who recently sold its agent-based electronic commerce system, Jango, to the Excite search engine company for \$35 million (and Excite’s share value had doubled since). Jango, an assistant shopper, simultaneously employs seven types of search engines to find the cheapest prices of books, software and other single-dimensional homogenous goods.

The current version can be found at <http://jango.com>. Other similar applications exist, including MIT's Firefly, which was recently acquired by Microsoft.

2.1.2 Industrial Applications

Industrial multi-agent systems are natural descendants of robot systems. Over the last three decades, ideas from co-operative and non-co-operative game theory have been used for solving co-ordination problems amongst independently controlled robots (for example, co-ordinated movement). Unlike robot systems, industrial MAS do not require that agents will have a physical existence, but instead use them as a metaphor for the inherently distributed process which underlies the application. Unlike search agents, which typically consists of a short code, agents in most industrial applications are computationally strong (typically each agent has its own CPU and I/O devices), capable of solving complex problems.

One of the first applications is ARCHON (Jennings et al. (1995)), a methodology and software platform for building multi-agents systems, used amongst other things in the process of particle accelerator control. Other applications include OASIS, a multi-agent system for air-traffic control, currently being used at Sydney airport (see Kinny et al. (1996)), and YAMS (Yet Another Manufacturing System), which uses Smith's contract nets (see Smith (1980) for more details) for the process of manufacturing control (see Parunak (1987)).

2.1.3 Business Applications

Nick Jennings's ADEPT (Advance Decision Environment for Process Task) views the management decision making process as a process of negotiating between various self-interested entities (see Jennings et al. (1996)). The system is successfully being used by British Telecom (BT), where agents represent the different departments, for example the legal department or the sales department, involved in the process of providing quotes and tariffs for consumers. As in most industrial applications, ADEPT agents are computationally strong. This is necessary because negotiation typically takes place over multi-dimensional services, where complex problem solving is required (for

example, at some stages of the negotiations an agent will need to solve a multi-dimensional constraints maximisation problem).

An important feature of ADEPT is that it allows for both external and internal negotiations (the latter could, in principle, be replaced with a centralised decision if negotiation fails). The system imposes only minimal restrictions on the negotiation process. On the one hand, this allows ADEPT to serve as the basis for many types of bargaining and negotiation scenarios, but on the other, it becomes difficult to obtain any general analytical results on the overall performance of the system.

2.1.4 Marketing Applications

Agent technology receives enormous attention from marketing companies. This technology has the potential to completely transform the way marketing is being practised: Agents can (and do) tailor advertising and web pages contents to the individual surfing. For example, based on the correlation with other consumers, who demonstrated similar interests, Firefly (www.ffly.com) recommends books and music CD's. A similar approach is taken by Yahoo!, with their advanced personalised search agent, MyYahoo!, (my.yahoo.com), which allows people to construct a news web page based on their preferences, while at the same time customising the commercial advertising on the page to the profile of the individual user and his or her behaviour.

Agents affect the marketing profession in several different ways, three of which are briefly described below (the interested reader is referred to Pazgal & Vulkan (1998) for a comprehensive survey on marketing applications). First, agents can be used by marketing companies, in order to create (and sell) consumers' profiles for a fee (for example, GolbalMedia uses an agent, known as "Rover", to construct special interest direct mailing lists). Second, host sites can interrogate visiting agents, similar to the questionnaires consumers are often asked to fill, thus collecting information and building up consumers' profiles. Finally, the increasing usage of automated agents is likely to change the way revenues from advertising on the WWW are distributed. Currently, most of the revenues comes from banner advertising, which is charged based on the

number of hits for the host site. But if most hits are from automated agents than this measurement is clearly no longer suitable.

2.1.5 E-Commerce

Currently, there are no large scale, fully automated, open-for-all, e-commerce systems. As explained earlier, in existing e-commerce applications users maintain control, and agents are only used as information gathering or information filtering devices. Still, a few small scale systems are already beginning to emerge. Two such systems are:

1. *Kasbah*, (see Chaves & Maes (1996)), a trading system, where users create simple buying and selling agents which then interact in an electronic marketplace, and
2. The American company Fastparts (see <http://www.fastparts.com/>) which provides a marketplace for electronic components, where users, human or automated, interact through a double auction mechanism.

Somewhat less known, but possibly more advanced than any of the systems mentioned above, is the Hebrew University “Popcorn” project (see <http://cs.huji.ac.il/~popcorn/>). This system views the whole of the Internet as one gigantic computer, where computations can be carried out in parallel in various locations over the network. A market-based mechanism is then used for the trade in CPU time.

3. Multi-Agent Systems

Economists are forever theorising about how they could make things more efficient. However, even if the theoretical foundations are well understood, there is always some degree of unpredictability when these models are put to practice. Even the well designed FCC auctions, perhaps the greatest practical success of implementation theory in the last decade, did not escape its fair share of signalling and price rigging allegations (it was reported that in one of the recent auctions, one of the bidders used the last few digits of amount bid to signal the telephone code of the area it was after. Although this was not coordinated between bidders in advance, the signal was

correctly interpreted by the other players, and the revenue generated was significantly lower than what was originally anticipated). At least with existing technology, computerized agents are unlikely either to be capable of generating such subtle signals or of interpreting them even if they are produced. It is therefore easy to see why the design of automated markets for software agents provides a perfect testbed for implementation theory. First, automated agents operate within the tight restrictions on the pre-specified protocol (mechanism) which controls the exchange of messages. The risk of signalling and price-rigging is therefore significantly reduced (if not completely eliminated). Second, unlike their human counterparts, software agents are time consistent entities which always choose optimally (given their computational ability, which is normally far greater than that of humans), much like game theoretical agents. In fact, game theory (and mechanism design) is much more suitable for automated agents than it is for people.

Automated agents may be less bounded than humans with respects to their computational abilities, but they are bounded by their knowledge bases, which raises new types of problems unfamiliar to traditional mechanism design. For example, an agent cannot initiate an auction mechanism, which could be in its best interests, unless it is specifically designed to do so. In general, the idea that agents *always* best-respond is questionable for automated agents. That is, it is extremely difficult to design automated agents which adapt to changing circumstances, but it is relatively easy to design agents who solve well define optimisation problems - quite the opposite case from humans⁵. Mechanism design for automated agents also takes into account the current state of supportive legislation, and issues related to security of transactions.

The discussion thus far suggests that existing mechanisms will need to be tailored for the specific needs of the MAS where they will be used (see, for example, Vulkan & Jennings (1997)). But this alone will not be sufficient, because many of these applications present us with new set of problems which have not yet been studied in sufficient detail by economists or game theorists. Consider, for example, the typical

⁵ As a matter on interest, this is a re-occurring problem for the science of Artificial Intelligence: AI systems can be taught to do what for humans is considered a hard task (e.g. expert systems), but fail to replicate simple everyday human behaviour, like the ability to understand human language.

scenario for one-to-one negotiations between automated agents where both agents face a strict deadline, which is private information, by which they *must* reach an agreement. The large literature on bargaining does not offer much insight to this problem. It would therefore be interesting to study the theoretical detail of a model which is inspired by this application. If we could find an analytical solution to some of these problems, like Rubinstein's unique subgame perfect equilibrium for the alternating-offer model (Rubinstein (1982)), then this is good news for the automated system, because negotiating agents can be pointed directly to it, hence bypassing the learning process. A similar approach can be taken for any protocol which yields a unique equilibrium. More generally, knowledge of the analytical solution(s) can be used by the designers of the system to improve its overall efficiency. It is therefore likely that in the near future we will see a large number of theoretical models inspired by agents and other e-commerce applications (like Monderer & Tennenholtz (1997a & b), or Sandholm (1998)).

A final difference is related to the distinction economic theory makes between mechanism design, where agents are assumed to behave optimally (i.e. attention is restricted to the set of Nash equilibria, or Subgame Perfect equilibria of the game), and models which take the structure of the game as given and study the optimal response of rational agents to the rules and to each other. In the current state of affairs of e-commerce, these two approaches are being explored simultaneously. In fact, most electronic commerce systems either design their own agents (like ADEPT), or offer the user the choice of using an agent design by them (like many of the on-line auctions sites). Until negotiation protocols become more standard, there seems to be little choice but to accept this dual approach. Still, it is useful to maintain, at least in our minds, a clear distinction between the design of agents, which optimise given the protocol, and the design of the protocol itself.

3.1 Evaluation criteria for MAS

Not surprising, the evaluation criteria for electronic systems are not very different from the general criteria used in implementation theory. For example, in an influential A.I. paper, Kraus, Wilkenfeld & Zlotkin (1995), list the following five points:

- *Autonomy*: the decision making process should not be centralised in any way, but fully distributed.
- *Promptness*: there should not be delay in reaching agreements (except, of course, for the case where delays are used for type-signalling).
- *Efficiency*: the negotiation outcome should be (Pareto) efficient.
- *Simplicity*: the protocol should minimise computation and communication resources available to the agents⁶.
- *Symmetry*: Whilst respecting roles, like buyer or seller, the protocol should not discriminate between agents.

Vulkan & Jennings (1997) consider an additional criterion based on the robustness of the MAS to the participation of non-optimising agents. This is particularly important in an open-for-all e-commerce systems, like on-line auction houses, where participating agents need only pass minimal compatibility checks. A system where a small number of non-optimising agents can significantly reduce its efficiency (say, by causing long delays) may be overall inferior to another system, which is robust, even if this comes at a cost of some small loss in efficiency when everyone behaves optimally. The relevance of this last point is further explained in section 3.5 below.

3.2 Contract Nets

Economic theory had long recognised that contracts which are *ex ante* efficient (in the sense that they maximise the utilities of the parties entering the agreement), does not necessarily have to be *ex post* efficient. When uncertainty is resolved, agents could find themselves in situations where they may have been better off not continuing with the original terms of the contract. Similarly, agents may find that the full commitment contract, which they choose not to enter because their outside option was *ex ante* preferred, would prove *ex post* superior. Instead, agents can sign contracts which are *contingent* on some probabilistically known, verifiable future event(s). It is a simple

⁶ Cost of computations can enter directly through agents' utilities. It is more difficult to do that with communication resources, which are a type of negative economic externality agents can enforce on each other.

exercise to show that there exist circumstances where contingent contracts outperform full commitment agreements (that is, there exist situations where two rational agents will not enter a full commitment contract but will enter some form of a contingency contract). The importance of these types of results for the efficiency of e-commerce systems with self-interested agents should therefore be clear (for example, scheduling agents or service providing agents will rarely enter a full agreement contract).

Unfortunately, there are several important reasons why contingency contracts are not suitable for multi-agent systems. First, it is difficult to allow unrestricted contingent contracts and in the same time control the complexity of the set of events on which the contract will be contingent upon. Second, it may be difficult, or even impossible, to verify whether some events actually happened. Instead, Sandholm and Lesser (1996), introduce *level commitment contracts*. In this framework agents specify penalties for unilateral decommitting from the agreement. Since level commitment contracts are not contingent on any future events, the problems of verification and complexity are avoided. New type of issues, however, needs to be considered, as agents may choose to delay decommitting, which otherwise would have been optimal, in the expectation that the other party will decommit first (there are double incentives to do that - avoiding paying the penalty, and gaining the penalty which is then paid to them by the other agent).

Using the tools of game-theory, by performing a full Nash equilibria analysis, Sandholm and Lesser show that in a fairly large set of cases (including the cases where agents decommitment decisions are taken strategically, as described above), their protocol outperform the full commitment protocol: First, there exist circumstances where level commitment contracts will be signed while full commitment contracts will not, while the reverse is never true. Second, compared to using the full commitment protocol, both agents are better off (in expected terms) using the level commitment contract protocol.

3.3 Auctions

On-line auctions account for a large share of Internet based trade. The number of items which are sold daily is estimated at hundreds of thousands. From the more well known auction forms, the English auction protocol is by far the most popular method used on-line (see, for example, <http://www.infohwy.com/nauck/vra19/PROTOCOL.htm>, for vintage records, <http://www.onsale.com>, for computers and electronics, <http://www.7cs.com>, for art, the Spanish fish market on-line auction (see Rodriguez et al. (1997), or <http://www.auctionline.com>; <http://www.interauction.com>; and <http://auction.eecs.umich.edu> which are general purpose). Some, like Klik-Klok (www.klik-klok.com), use the Dutch auction protocol and others, like Insurance Auto Auctions (www.iaai.com), use first-price sealed bid auction protocol. Many of the sites which uses English auction allow users to bid indirectly using agents (known as “bidding elves”) which keep increasing their bid until either they win the auction, or until they reach a price pre-specified by the user. As long as bidders report truthfully their reservation prices to their agents, then the auction will be profit maximising for sellers. Auction theory provides us with an understanding of the conditions under which the various auction forms are optimal. However, Internet auctions are different in some important respects: Because the items sold are mostly cheap (for example computer parts and TV sets), the set-up costs and participation costs are relatively low. Moreover, both sellers and buyers can choose between the different on-line auction houses. In a recent paper, Monderer & Tennenholtz (1997a) formalise some of these important differences, and in particular focus on the facts that (1) sellers compete in auction forms rather than in prices, and that (2) bidders are mostly risk-seeking, rather than risk averse. They then prove that, in the presence of risk-seeking agents, a seller is better off using a third price auction protocol. In general, Monderer & Tennenholtz specify conditions under which a k price auction is revenue superior to a $k-1$ price auction.

Auctions provide a quick and efficient way of resolving one-to-many negotiating situations (see, for example, Binmore (1985), which shows that sequential negotiations between rational agents leads to the same outcome of an auction. If time is discounted or if communications are costly, then it is easy to see why an auction is more efficient). It is therefore not surprising that many MAS use auctions as internal service and resource allocation devices. An internal auction is generally similar to any other an on-line auction. An important difference is that the auction has to be initiated by any agent

which has to negotiate with several other agents, each of which can potentially provide the required service. The MAS is therefore only efficient as far as its agents are able to recognise such situations, and their ability to initiate efficient mechanisms in appropriate circumstances. This may seem easy for human agents (one can always suggest to an individual or a firm the option of using an auction), but is in fact difficult for self-interested automated agents.

Vulkan & Jennings (1997) provide a formal analysis of the efficiency implications of using English auctions for internal resource allocation in MAS. The analysis is carried out in the context of the ADEPT system (see section 2). Vulkan & Jennings (1997) show formally how a *pre-auction protocol* can be employed by the bidders to restore the efficiency of the system in situations where the service-seeking agent fails, for any reason, to initiate the (efficient) auction protocol. A pre-auction is an auction which is held by the bidders alone. The winner of this auction is the only agent who will then agree to negotiate with the service-seeking agent. Vulkan & Jennings (1997) show that, under some conditions, such a mechanism can be incentive compatible and can lead to an outcome similar to that of an efficient auctions. Pre-auctions, or in general, efficient mechanisms which generate the same outcome as auctions, but which can be initiated by agents from either side of the negotiating sides, may prove an important feature of mechanism design for computerised agents.

3.4 Automated Negotiations

The classical problem of how to split a (possible shrinking) pie between two self-interested agents lies in the heart of many MAS applications. It arises, for example, when scheduling agents bargain for preferred meeting times for their users, or when agents negotiate the final price of a good, or when agents representing different departments bargain over the small prints of a service which they provide jointly (as in ADEPT, see Jennings et al. (1996)). In an automated system the designers choose the *negotiation protocol*, which will then be used by the negotiating agents. For example, protocols may be chosen where agents are allowed to make offers and counter-offers. In economic theory attention is normally focused at mechanism which resembles how people behave, but this is no longer an advantage for automated agents. For example,

the protocol may interfere with the negotiations if both agents do not make any concessions (as in earlier versions of ADEPT). Of course, agents operating under such a protocol may optimally respond to this by deliberately delaying making any concessions in anticipation that their opponent will concede first in order keep negotiating.

Perhaps the most popular descriptive model of bargaining studied by economists is Rubinstein's alternating offers game (Rubinstein (1982)). Since the model has a unique solution where agents agree a split immediately, it seems particularly attractive for automated bargaining solutions, because it satisfies all five efficiency criteria specified above (see, for example, Rosenchein & Zlotkin (1994), and Kraus, Wilkenfeld & Zlotkin (1995)). However, there are two major difficulties with the above argument: First, it is known that the outcomes of the model are very sensitive to many of the assumptions, like the specific form of exponential discount rates, or the fact that the good must be infinitely divisible. This means that a MAS which uses this protocol, is only efficient as far as these assumptions are met. For example, under linear time discounting (i.e. a fixed cost is incurred for every round of offers), the results of the model change dramatically, and the first mover either gets the whole surplus or most of it (depending on the ratio between the fixed costs for the two players, see Osborne and Rubinstein (1990)). Since automated agents can carry out millions of rounds of offers and counter-offers every second, exponential time discounting may not be the most appropriate method for modelling players' aversion to delay. Second, the model has an additional disadvantage for automated interactions in that it is complicated for the case of two-sided uncertainty. In general, agents may choose to deliberately delay agreements in order to signal their type. Moreover, the problem of multiple-equilibria rises in a particularly aggravated form in such situations. Equilibria exist where the true type of agents is not revealed, and when true types are revealed after long delays (in general the length of the delay will depend on the number of possible types). The usefulness of the model can be questioned when it allows for such different outcomes. Since a large amount of uncertainty is almost always present in bargaining between automated agents, it is difficult to see how one can get around this problem.

The insights from bargaining theory can still prove useful through the study new types of bargaining models, inspired by the various applications of automated negotiations. Any information about the type of the negotiating agent (such as discounting or deadline) must be assumed to be private. At least some of the success of the alternating offers model is due to the fact that it resembles how people bargain, but this is no longer an advantage for automated interactions. The temporal monopoly assumption, which underlies Rubinstein's model, can be replaced with a mechanism where agents can make or accept offers at any stage of the negotiations. Moreover, in most of the MAS applications described in this paper agents face deadlines (i.e. a time after which there is a sharp drop in utility), which are far more important to the final agreement than individual time discounting. Taking on board the last two points, bargaining between automated agents may be better modelled as a variant of the war of attrition, instead of alternating offers. Of course, this model may prove too difficult to analyse analytically, but even a partial understanding of what is optimal behaviour in such circumstances can prove useful for the design of negotiating agents and negotiation protocols.

3.5 Knowledge Representation

This purpose of this section is to illustrate how the theory of mechanism design can be used towards constructing efficient protocols for multi-agent systems, when it cannot be assumed that the agents know in advance the full details of the environment. The main finding is that, even in relatively simple situations, mechanisms which are not dominant-solvable may prove inefficient because of difficulties with the representation of knowledge about reasoning about the behaviour of other players.

Consider the following co-ordination game represented in normal form in figure 2 below. The game have two Nash Equilibria, {Up, Left}, and {Down, Right}. Consider now the task of writing a computer program which may have to play this (or a similar) game.

	Left	Right
--	------	-------

	Left	Right
Up	2,2	0,0
Down	0,0	2,2

Figure 2: A co-ordination game

It is easy to see that there does not exist a programme which guarantees a payoff of 2, because the outcome will always depend on the action of the other players. Moreover, trying to reason about the behaviour of the opponents in this situation is not likely to be useful because there are no focal points⁷. An agent reasoning about its opponent can run into an infinite loop, because the reasoning of agent A on B's behaviour depends on what B reason on A, which depends on what A reason about B reasoning about A, ad infinitum. Running into an infinite loop is not only bad for the agent, but could also negatively affect the overall performance of the MAS.

In contrast, consider the task of writing a programme which plays the prisoner's dilemma (figure 3 below):

	Left	Right
Up	2,2	0,5
Down	5,0	1,1

Figure 3: The Prisoner's dilemma

It is now possible to programme an agent to never play a strategy which is dominated, and to never expect its opponent(s) to do so either. First, there is some hope that it is possible to represent the statements above in terms of a computer code. Second, an agent equipped with such reasoning can be expected to reach the equilibrium of any game which is (finite) dominantly solvable, like the prisoners' dilemma.

The above example illustrate the following point: Given the complexity of knowledge

⁷ In general, even situations which have a unique focal point are not easy to solve, because we do not yet have a uniform theory of what is focal. Still, some types of focality arguments, like payoff dominance, are relatively easy to represent.

representation induced by counter-speculations, mechanisms which rely on dominant strategies are clearly desirable. Moreover, it should be easy to see that these mechanisms are, potentially, more robust against the participation of a small number of non-optimising agents (because what is optimal for an optimising agent does not depend on the behaviour of its opponent). This point can hardly be stressed enough: A MAS which relies on a protocol which is *not* dominant solvable, *cannot*, in general, *be expected to induce optimal behaviour* from agents with limited knowledge base. It is therefore extremely important that any effort will be made to use dominant solvable mechanisms. Implementation theory provides us with general methods of converting a given mechanism to an equivalent dominant-solvable mechanism. Such methods could therefore be put to use in order to ensure the robustness of multi-agent systems.

3.6 Strategic Ignorance

Mobile agents which create copies of themselves over the Internet, can become vulnerable to hostile host sites. The hostile host can either directly observe their code, or could execute a copy of the agent in order to find out details about its behaviour. This raises obvious difficulties for the designers of negotiating agents. While advances in cryptography theory are expected to provide a solution to this problem, Vulkan (1998) shows that there exist circumstances where it becomes in the best interest of an agent to maintain the option of revealing its code. An automated agent claiming to not be able to negotiate further (for example, not allowed to reduce the price with which it is selling a given object), can credibly prove such a claim by revealing its code. This is rarely the case with most other types of negotiations (even when firms use intermediaries to negotiate on their behalf), where such claims are often being made untruthfully exactly because they cannot be verified.

More formally, the intuition for this result is the following: by revealing its code, an agent can credibly signal that it is unable to distinguish between certain states, or forms of behaviour, thus changing the structure of the underlying game. This can be beneficial for the agent if it is better off in the outcome of the new game compared to the original situation. The most obvious example is a selling agent which credibly shows that it is unable to bargain, and can only accept a given price for the good it is

selling. In the equilibrium of this take-it-or-leave-it game, the buyer either takes it or leave it, an outcome which may very well be preferred by the seller to any form of direct negotiations. Vulkan (1998) shows that this intuition can be generalised to different game forms, and may even result in outcomes which are very different from that of the original game (in the simple example above, the seller is able to pick its preferred equilibrium by not bargaining. In contrast, Vulkan (1998) shows that under some condition it becomes possible for agents to co-ordinate on an outcome which is not at all supported by a Nash equilibrium of the original scenario). These type of results are robust to changes in technology: as long as users have incentives to create agents who credibly reveal their code, they will continue to do so.

4. Economic Implications of Trading using Agents

4.1 Search agents

The process of electronic commerce on the World Wide Web can be thought of as follows: Suppliers of a particular product announce on the Web the price at which they are willing to sell it. A consumer wishing to buy a product can then search the Web in one of two ways: The consumer uses a standard search engine, types in the name of the product, and gets a list of web-sites that supply the product. Given that these suppliers may come from all over the world, the list of web-sites can be enormous. The consumer then has to visit the web-sites individually to find the price. This can be very time-consuming, and realistically the consumer will only be able to visit a limited number of sites. The alternative is to employ a search agent. Here the consumer types in the name of the product, and the agent can *potentially* visit all sites, interrogate them to find the price, and returns the lowest price - or perhaps a range of prices. The consumer then decides whether and which to buy.

From the description above, it should be clear that a search agent significantly lower search costs - and effectively drive them to zero. Given this technology, the sorts of questions economists should be investigating are:

- (i) What types of products can/will be traded using search agents?
- (ii) What will be the impact of the use of such search agents on product prices?
- (iii) What will be the impact of such search agents on the range and quality of products on offer to consumers?

At first sight one would think that this would make markets more competitive and lower prices. However, the following three points poises some doubts on this argument:

- (i) Suppliers can block their sites from being interrogated by search agents. Indeed many of the companies which initially allowed BargainFinder to use their site, blocked such access in later stages of the experiment, while others who did not allow access in the first stages, later reversed this decision.
- (ii) Agent technology lowers the search costs not only of consumers but also of suppliers who wish to find out what prices their rivals are charging. This makes it easier for firms to operate trigger-price strategies which may enable them to sustain high prices. This also makes it more difficult for sellers to secretly undercut each other.
- (iii) Some consumers will not use search agents: This is either because they do not have access to this technology, or because they do not know how to use it, or because they are reluctant to delegate decisions (and financial decisions in particular) to, what is essentially, an artificial agent. The relatively small numbers of users of the more sophisticated search engines (such as AltaVista Advanced search, or MyYahoo!), despite their accessibility and easy-to-use interfaces, suggests that the last point could prove, at least initially, to be significant.

One implication of this is that even though consumers can costlessly determine the prices of those suppliers who choose to allow their site to be interrogated by search agents, they have to decide whether, in addition, they should also engage in some costly search of the suppliers who do not allow their site to be interrogated by search agents. Obviously one factor that will govern this choice will be the number of firms in each group. This generates a clear network externality whereby the decision of a firm

whether or not to block access of search agents, indirectly affect the decisions of consumers about which search technology to use.

To understand the implications of the existence of search agents we therefore need to understand fully the underlying incentive structure of sellers and buyers given the above assumptions. Specifically, we want to examine the equilibria of such models to see (i) how many suppliers choose to allow their site to be interrogated by search agents; (ii) the prices set by suppliers who permit interrogation, and those who block interrogation; (iii) the search and purchase strategies of consumers.

While this model appears simple no such model of pricing behaviour currently exists. Models with costly search typically assume exogenous search costs (as in Chung & Lee (92), Hwang (93, 95), Li, Mckelvey & Page (87), for example), whereas what gives this model its interest and bite is the fact that suppliers can determine which technology others must use to ascertain their price. The endogenous information acquisition approach (as in Vives (88), or Hurkens & Vulkan (96, 97)) may be more complicated, but is more suitable for the setting described above. The effects of costly information acquisition on economic equilibria were further explored by the literature on bounded rationality (e.g. Abreu & Rubinstein (88), Rubinstein (93), Dow (93)). Dow's model investigate the optimal response of firms to consumers with limited search power. Whether Dow's intuitions will remain useful in a setting where consumers can choose their search power remains to be seen.

First and foremost it is important to study the equilibria for the case of homogenous goods. Currently, these types of goods account for the bulk of the trade on the Internet (because otherwise it is more difficult to define a search pattern). However, with newer, more sophisticated agents being developed, there is no reason why other types of goods and services will not be traded in the near future. The importance of a sound theoretical framework is clear, in order to study and predict the effects of using agents on the markets for such goods. As suggested by the models of Hurkens & Vulkan (1996), and Fershtman & Kalai (1993), in markets for goods with a degree of product differentiation, a possible outcome is that sellers endogenously specialise in the different attributes of the good or service in question, hence securing a niche for

themselves. More formally, for a given demand and supply schedules sellers could either diversify their offerings and split the market, or they could specialise. Although both types are payoff equivalent, only the latter is robust against an increase in consumers' search costs⁸.

4.2 Other Implications of E-commerce

The framework discussed so far, takes as given the participation decisions of sellers and buyers. But e-commerce is also likely to have structural effects, at least for some types of businesses: In the immediate future, it has important implications for employment of existing intermediary agencies such as travel agents. In the long run, agent technology has the potential to effect the future of many different trading institutes, because of its significant reduction in set-up costs for new businesses (this is particularly important for the retail industry). Because electronic shopping assistance search and buy in virtual shops, we may experience, in the long run, some effects on real-estate prices, through reduced demand for physical outlets.

Agents can be used to replicate the function of organisations whose main business is to serve as intermediaries: the main reason I visit my travel agent, is because he can find out the most suitable deal for me relatively quickly (especially if he knows me). But if my software agent can search, potentially a very large set of options, and match directly to my preferences, then this trip down the high street may no longer be necessary. The strength of human agencies lies in its ability to provide exceptional access or superior information. These services will need to adapt to the new rules of the game or be eliminated.

The above claims are not altogether academic as experience with Internet travel agencies shows: The two leading Internet travel agencies, Microsoft's Expedia (www.expedia.msn.com) and Saber's Travelocity (www.travelocity.com), are capturing large segments of this market, with Expedia's sales expected around \$110M in 1997, and

⁸ Fershtman & Kalai prove this result in the context of complexity constraints, while Hurkens & Vulkan prove a similar result in the context of costly information gathering for firms entering a new

Travelocity's revenue exceeding \$95M in 1996 (source: Warner (1998)). The problems faced by conventional travel agencies and other intermediaries are real and urgent. Unless existing agencies can find a way to utilise this new technology while maintaining superior matching skills, their future is uncertain.

5. User-Agent Interaction

The design of interfaces for automated agents raises interesting issues from the point of view of economic theory. The main difficulty lies with the fact that although people are rarely conscious about their preferences, software agents cannot be fully autonomous unless they know the utilities of their users. A second major difficulty relates to trust: if agents are designed by organisations which might profit from particular forms of behaviours, then users may be reluctant to reveal all relevant information to their agents.

The interactions between users and agents are currently being investigated by several commercial organisations. In particular, Hewlett-Packard's JESTER experiments, where users interact in controlled environments, and under various conditions, so that data can be collected on the type and level of communications between users and agents, and the associated degree of trust, as measured by users' willingness to delegate tasks to their agents (see Preist & Van Tol (1998), for more details).

Along the same lines, Anderson introduce their experiment, LifestyleFinder agent, Waldo the Web Wizard (see <http://bf.cstar.ac.com/lifestyle>), a marketing oriented application where inferences about the user are drawn from choices related to lifestyle. The user is faced with series of simple multiple choice questions (for example, the user is asked about the type of programmes he or she likes to watch on TV). LifestyleFinder is an example of an interface which very quickly builds a user profile, which can then be used as the basis of independent decision making.

market. The intuitions from either results, however, hold for the case where costs are imposed on buyers' search.

5.1 Usage of Agents

To enable automated agents to internally represent their user's utilities, designers should first consider the type of tasks and the length of time that the agent will be used for. If an agent is going to be used repeatedly for a set of relatively similar tasks (for example, shopping or scheduling), then it is possible to draw inferences about the user's *direct* utility from observing his or her behaviour (i.e. from observing the user's *indirect* utility). Economic and econometric theories can be used in constructing the learning algorithms for the agent. First, the agent can be provided with a prior over the set of possible utility functions for its user. Keeping in mind that agents only trade in goods and services with a well defined list of attributes, this set may not be that large. As more data becomes available, the agent can simultaneously estimate the parameters for each utility form and compare the likelihood of each model.

Of course, if the agent is only used a small number of times, then this method is no longer appropriate. Instead, the agent can prompt its user to reveal as much information through a series of pre-designed questions. As Anderson's experiments with LifestyleFinder show, if the agent start with a reasonable set of utility forms (for example based on consumer profiles), even a small number of questions can quickly, and with high probability, converge to the right type.

Next, the type of information the agent will need to obtain from the user will depend on whether it is a task specific or a general purpose agent. A general purpose agent will need to find out general features of the preferences of its user, like attitude to risk, preferred level of service, the user's budget constraint, etc. In contrast, a type specific agent, for example an agent which buys particular types of services on behalf of the user, will need to estimate its user's specific trade-offs between the various attributes and his or her reservation prices (if those exist). For example, the agent will need to know its user's trade-off between price and quality of service.

5.2 Trust and Truth Telling

The issue of trust between users and software agents is clearly important for the future success of this technology. Economic theory suggests that this will, amongst other things, depend on the incentives of whoever wrote the agent. For example, some on-line auction sites offer the option of using one of their agents. This may be innocent enough, but if the auction house gets a percentage of any trade which takes place on that site, then there is a clear incentive to design agents such that users are more likely to trade. Or more generally, users might be suspicious about using agents written by large corporates, which may have their own agenda for the future of electronic commerce.

The importance of issues of trust was highlighted in some of the early experiments with electronic marketplaces (for example in MIT's *Kasbah* experiments), where it was reported that some users consistently lied to their agents about their reservation prices, presumably because they expected that will cause their agent to be "tougher", which may then lead to a better outcome for them. What game-theory teaches us, by means of the revelation principle, is that, as long as the interactions between the user and her agent are kept secret, then the user is not worst off, and possibly better off, telling the truth about her type. That is, if lying is profitable, then the agent can be trusted to do all the pretending necessary. It is of course a challenge for the designers of agent interfaces to make this clear to users who are less familiar with such theorems.

5.3 Agents as Commitment Devices

There is a literature in economics which studies the implication of delegation to intermediaries of the responsibility of negotiating on behalf of firms. Notably, Fershtman & Judd (1987), and Fershtman, Judd & Kalai (1991), who's notion of an agent is quite similar to that of a software negotiating agent. The intuitions gained from their models may therefore prove useful in our understanding of automated negotiation. Fershtman & Judd's main result is that when the choice of the agent becomes observable, it is possible that new agreements, which are not supported by an equilibrium of direct negotiations, could become an outcome of the agent-choosing game. In a recent paper, Vulkan (1998) shows that, in the presence of some uncertainty, users will sometimes prefer to choose "ignorance" agents so to increase

their commitment power. As explained in Section 3, software agents are unique in that they can credibly reveal their “negotiation instructions” to their opponents. Utility maximising users might therefore pre-commit by choosing a deliberately restricted negotiating strategy - through their choice of an agent. If a user intend to use its agent in such a way, this have immediate implications for the design of agent interfaces. At a minimum, agents could be design in a way which enables them to play such roles.

6. Conclusion

Software agents can now gather and filter through the vast amounts of information available on the Internet. Existing technology now makes it possible for agents to schedule meetings and negotiate prices and agreements on behalf of their users. How these developments will effect markets in still largely unknown. This paper discusses the types of research which could be undertaken by economists in order to have a better understanding of this question.

Because agents search for a pre-specified pattern, like price, they cannot be influenced by other features of the product, and the scope for product differentiation is therefore significantly reduced. Firms which allow software agents into their sites might find that they are forced into a Bertrand-like competition scenario. Are we likely to see prices derived down to costs? As explained in section 4, this is not at all clear. As experience from the BargainFinder experiment shows, firms can still manage to product differentiate through other means, like bundling.

Multi-agent systems are studied and designed by researches from a field of computer science known as distributed artificial intelligence (DAI). It is an interesting fact that artificial intelligence (AI) and economics have had many overlapping interests over the years. John Von Neuman’s pioneering work had laid the foundations for modern AI as well as modern game theory. Along the same lines, Herbert Simon’s work on rationality and bounded rationality had greatly influenced researches in both fields. It is therefore not surprising that we find ourselves these days in a situation where

researchers from both fields work together in pursuit of what may become one of the more important technological changes of modern life.

In the short run, computer science is likely to benefit more from this co-operation, because economic wisdom on the efficiency of systems consisting of self-interested agents can be almost directly applied to multi-agent systems. In addition, the economic methodology which stresses the importance of looking for the underlying incentives of participants can provide important insights into automated negotiations and electronic trade. But in the longer run, economists have also much to gain from this joint adventure. By focusing on what is essentially an application of economic theory and mechanism design to automated environments, we can learn about the usefulness of our theories and intuitions.

References

Abreu, D. and A. Rubinstein (1988), “The Structure of Nash Equilibrium in Repeated Games with Finite Automata”, *Econometrica*, pp. 1259-1281.

Binmore K. (1985), “Bargaining and Coalitions” in Alvin E. Roth (ed.) *Game Theoretic Models of Bargaining*, Cambridge University Press.

Chang, C. H. and C. W. Lee (1992), “Information Acquisition as Business Strategy”, *Southern Economic Journal*, no. 58, pp. 750-761.

Chavez A., and P. Maes (1996), “Kasbah: AN agent marketplace for buying and selling goods”, *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Systems*, London, UK.

Dow, J. (1991), “Search Decisions with Limited Memory”, *Review of Economic Studies*, **58**, 1-14.

Fershtman, C. and K.L Judd (1987), “Equilibrium Incentives in Oligopoly”, *American Economic Review*, **77**(5), 927-940.

Fershtman, C., Judd, K.L., and E. Kalai (1991), “Observable Contracts: Strategy Delegation and Cooperation”, *International Economic Review*, **32**(3), 551-559.

Fershtman, C. and E. Kalai (1993), “Complexity Considerations and Market Behaviour”, *RAND Journal of Economics*, **24**, 224-235.

Gilbert D., Apaecio M., Atkinson B., Beady S., Ciccarino J., Grosf B., O’conner P., Osisek D., Pritko S., Spagna R., and L. Wilson (1995), “IBM Intelligent agent strategy”, IBM Corporation.

Guilfoyle C., and E. Warner (1994), "Intelligent Agents: The new Revolution in Software", *Ovum Ltd.* London.

Hwang, H. (1993), "Optimal Information Acquisition for Heterogeneous Duopoly Firms", *Journal of Economic Theory*, no. 59, pp. 385-402.

Hwang, H. (1995), "Information Acquisition and Relative Efficiency of Competitive, Oligopoly, and Monopoly Markets", *International Economic Review*, no. 36, pp. 325-340.

Huberman B.A, and R.M Lukose (1997), "Social Dilemmas and Internet Congestion", *Science*, Vol. 277, 535-537.

Hurkens S. and Vulkan N. (1996), "Information Acquisition and Entry", University College London discussion paper 96-23

Hurkens S. and Vulkan N. (1997), "Free entry does not imply zero profits", mimeo, Bristol University, and Universitat Pompeu Fabra.

Jennings N.R., Corera J. M., and I. Laresgoiti (1995), "Developing industrial multi-agent systems", *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*, 423-430.

Jennings N.R., Farain P., Johnson M.J., Norman T.J., O'Brein P. and Wiegand M.E. (1996), "Agent-based business process management", *International Journal of Co-operative Information Systems*, 5(2-3): 105-130.

Jennings N.R., and M. Wooldridge (1997), "Applications of Intelligent Agents", <http://www.springer.de/comp/special/jennings.pdf>

Kraus S., Wilkenfeld J., and G. Zlotkin (1995) "Multiagent Negotiation Under Time Constraints", *Artificial Intelligence Journal*, 75(2), pp. 297-345.

Li, L., Mckelvey, R.D., and T. Page (1987), "Optimal Research for Cournot Oligopolists", *Journal of Economic Theory*, no. 42, pp. 140-166.

Mackie-Mason J.K., H.R. Varian (1995), "Pricing the Internet", In Kahin B., and Keller J. (Eds), *Public Access to the Internet*. MIT Press, Cambridge, MA.

Monderer D., and M. Tennenholtz (1997a), "Internet Auctions", mimeo, Technion - Israel Institute of Technology, Haifa, Israel.

Monderer D., and M. Tennenholtz (1997b), "Distributed games", mimeo, Technion - Israel Institute of Technology, Haifa, Israel.

Osborne, M. J. and A. Rubinstein (1990), "*Bargaining and Markets*", San Diego: Academic Press.

Parunak H.V.D. (1987), "Manufacturing experience with the contract net", In M. N. Huhns (Ed.) *Distributed AI*. Morgan Kaufmann.

Pazgal A., and N. Vulkan (1998), "Have you agent call mine: Software agents, the Internet, and marketing", mimeo, John M. Olin School of Business, Washington University, St. Louis.

Preist C., and M. Van Tol (1998), "Economic Agents for Automated Trading", Agent Technology Group Discussion Paper, Hewlett-Packard Labs, Bristol, UK.

Quarterman, J. (1995), *Matrix News*, Austin, TX, <http://www.mids.org>

Sandholm T.W., and V. R. Lesser (1996), "Advantages of a level commitment contracting protocol", in *Proceedings of the National Conference on Artificial Intelligence*, Portland, Oregon, 126-133.

Sandholm T.W. (1998), "Agents in Electronic Commerce: Component Technologies for Automated Negotiation and Coalition Formation", mimeo, Washington University, St. Louis.

Sierra C., Faratin P., and Jennings N.R. (1997), "A service-oriented negotiation model between autonomous agents", in the proceedings of Modelling Autonomous Agents in a Multi Agent World (MAAMAW) 1997.

Smith R.G. (1980), "The Contract Net Protocol: High-level communication and control in a distributed problem solver", *IEEE Transactions on Computers*, **C-29**(12), 1104-1113.

Rodríguez, Juan A., Noriega, P., Sierra, C., and Padget, J. (1997), "FM96.5 A Java-based Electronic Auction House", Paper prepared for the Second International Conference on *The Practical Application of Intelligent Agents and Multi-Agent Technology*: PAAM'97.

Rosenschein J., and Zlotkin G. (1994), *Rules of Encounter*, The MIT press, Cambridge, USA.

Rubinstein A. (1982), "Perfect Equilibrium in a Bargaining Model", *Econometrica*, **50**(1), 97-109.

Rubinstein, A. (1993), "On Price Recognition and Computation Complexity in a Monopolistic Model", *Journal of Political Economy*, no. 31, pp. 473-484.

Varian R. H. (1995), "Economic Mechanism Design for Computerised Agents", paper delivered at the USENIX Workshop on Electronic Commerce, July 1995, New York.

Vulkan N., and N.R. Jennings (1998), "Efficient Mechanisms for the Supply of Services in Multi-Agent Environments", mimeo, Economics Department, Bristol

University and the department of Electronic Engineering, Queen Mary and Westfield College.

Vulkan N. (1998), "A Note on the Value on Information in Automated Interactions", mimeo, Economics Department, University of Bristol.

Vives, X. (1988), "Aggregation of Information in Large Cournot Markets", *Econometrica*, no. 56, pp. 851-876.