Beyond Intelligent Buyers and Sellers: Introducing Manager Agents

Margaret H. Belknap
Department of Systems Engineering
University of Pennsylvania
Philadelphia, PA 19104-6315
ph: 215-848-0026 fax: 215-573-2065
belknap@seas.upenn.edu

Patrick T. Harker
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6364
ph: 215-898-4715 fax: 215-573-5001
harker@wharton.upenn.edu

Beyond Intelligent Buyers and Sellers: Introducing Manager Agents

ABSTRACT

This paper establishes the efficacy of an intelligent manager agent in the electronic commerce domain. Currently, intelligent agents in electronic commerce perform tasks as buyers, sellers, and search agents. We believe that as this form of commerce grows, it may become necessary for these agents to organize themselves into clusters headed by a manager agent. The manager agent may perform tasks similar to those of a human manager. By introducing a manager agent we increase efficiency in the market place in two ways. Agents that are best suited to service them serve more customers and the agents' revenue is increased. We use queuing theory to model a two-server type, two-client type system and offer a numerical example to illustrate our conclusions. In our simple model, the manager only performs the tasks of seeking and assigning work. However, we discuss a number of other tasks that a manager agent would be able to perform.

Keywords: electronic commerce, manager agent, intelligent agent, queuing, disintermediation, market intermediary.

1. Introduction

Electronic commerce is so large and is growing at such a rapid pace that many experts believe any reports on its size or rate of growth are obsolete by the time they are estimated. The Secretariat on Electronic Commerce at the U.S. Department of Commerce (USDC, 1998) and the Organisation for Economic Co-operation and Development (OECD, 1999) provide comprehensive reports on the status, potential, and impact of the emerging electronic market place. Both sources point out the difficulties of measurement in this market for several key reasons that include: the speed at which it is growing, the fact that it is still in its embryonic stage, and the reluctance of participants to share proprietary information. In spite of these difficulties, we can still develop models and apply basic theories to obtain important insights on issues related to electronic commerce. Shapiro and Varian's (1999) book on the durability of economic principles in the "network economy" stands as an example of such insights. Similarly, we propose that durable principles and theory from operations research and management science can be applied to gain important insights regarding this new form of commerce. In this paper, we apply queuing theory to model a key player, the intelligent agent in the world of electronic commerce.

Intelligent agents are used in a broad range of applications from the help wizard on desktop PCs, to the search agent on the Internet, to a system of autonomous agents for process control in industry, to agents bidding for resources in complex systems. The latter application, called market-based control, is a technique to allocate resources based upon the fundamental concepts of market mechanisms (Clearwater, 1996). This approach creates an artificial economy, usually an auction, where intelligent agents determine prices for resources through trade and, thus, their allocation. Examples include allocation of CPU time in computer systems (Chavez *et.*

al, 1997) supply chain management (Swaminathan *et. al*, 1998), factory scheduling (Baker, 1996), and conditioned air flow control in an office building (Clearwater and Huberman, 1994).

In electronic commerce, the intelligent agent is a hybrid of the search agent on the Internet and the buyers and sellers in the market-based control auctions. An intelligent agent can seek out opportunities for profit, participate in auctions, win bids and execute the order that they bid for. While intelligent agent development has been largely left to the computer and information scientist, the tasks they perform are familiar to the operations research/management science analyst. Their functions entail a multitude of tasks that include, for example: marketplace search, bid price calculation, production scheduling, inventory management, job execution, and delivery of the promised order.

The growth of research on the use and development of intelligent agent systems reflects the growing importance of this technology in practice. Microsoft Research was recently a co-sponsor of a five-day conference*, The Third International Conference on Autonomous Agents (Agent '99)*, entirely devoted to intelligent agent research. In addition to exploring the diverse domains in which intelligent agents can be applied, research issues included agent design, protocol, language structure, and agent similarities to humans. In electronic commerce, these investigations have been extensions of research done on Internet search agents and market-based controls. Topics include such areas as reducing buyer search costs (Bakos, 1997), and are primarily focused on how to calculate optimal bids (Sandholm, 1993) and how to structure the rules for auctions (Parkes, 1999). However, no one has focused on modeling optimal behavior for *work-seeking agents* that must allocate their time between seeking and performing work whether they are in electronic commerce or in a computer network system. An example of a working seeking agent in electronic commerce is one who must allocate its time between seeking opportunities to bid and calculating optimal bids. This paper is a first attempt at modeling work-seeking agents in general, and we choose electronic commerce as our domain.

At first glance, some may say that this problem of agents seeking and performing tasks is a non-issue due to the ability of computers to multi-task. However, from the perspective of queuing theory, multi-tasking does not do away with the fundamental problem that agents have a finite capacity in terms of computing power. Thus, there will always exist the problem of allocating this finite power between tasks.

We contend that as the marketplace grows, those agents that must both seek and perform work may become overwhelmed to the point where they can no longer participate efficiently in all of these tasks. It may be that as the market expands, these intelligent agents will find it necessary to organize themselves into clusters headed by a "manager agent". That is, such intelligent agents may require the services of a middle manager to represent them in the marketplace. In this sense, the electronic marketplace comes to resemble real markets where middle managers or "manager agents" (e.g., the partner in a consulting firm) seeks work for a group of agents that actually perform the tasks (e.g., the actual consultants) as well as providing oversight for these agents, certifying the quality of the work, and numerous other supervisory tasks.

By modeling the individual work-seeking agents, we can explore the potential impact of a manager agent's presence in the electronic marketplace. For our purpose, we will only consider the agents' seeking, accepting, and serving activities. In the absence of a manager, an agent must divide its time between seeking clients in the marketplace and servicing them because an agent cannot seek and serve clients simultaneously (as an example, think of the independent consultant). Therefore, clients that arrive to the marketplace while the agent is busy servicing another client either choose another agent or are lost. We introduce a manager agent and argue that it will increase market efficiency. The manager agent will assume the roles of seeking clients and assigning them to agents. As a result, the introduction of the manager agent will increase revenue for the agents and decrease the number of clients that are lost. Hence, both the agents and the clients are better off.

It is important to distinguish our advocacy for a middle manager agent from the current debate regarding disintermediation. Disintermediation is a term that was coined with regard to the financial services industry in the late 1960's to describe the trend for small investors to invest directly in financial instruments such as money market funds rather than through the traditional intermediary, a bank savings account (Gellman, 1996). In electronic commerce, many (e.g. Duncan, 1998; Wilder, 1998) believe that the application of information technologies has led to disintermediation, meaning the elimination of the middleman, the retailer. Malone *et. al* (1989) briefly describe an evolution from computer-aided buying and selling to electronic markets. They predict a decrease in company vertical integration. Arguing for evolution, not extinction, others (Fox, 1999; OECD, 1999; USDC, 1998) contend that electronic commerce will change the way the middleman does business. Alsop (1999) advocates looking for ways to use information technology to enhance the middleman's role. He also points to the surge in venture capital interest in web-based companies as evidence of a widely-held belief in Amazon.com as a business model of the future.

In this paper, we are not discussing the middleman, the retailer. We are introducing the possibility of a middle agent that would act as a manager for many intelligent agents. Our middle manager agent is also significantly different than an electronic intermediary. Electronic intermediaries more closely resemble the functions of a middleman, a retailer. Bailey and Bakos (1997) discuss two examples of electronic intermediaries and their reasons for success and failure based upon market forces. Their intermediaries performed a search and match service for customers. Our proposal is that a middle agent could allocate tasks to agents based upon their expertise. Further, we show why this is more efficient for customers and servers with a simple application of queuing theory.

In the next section of this paper we discuss modeling attempts and research focus as it pertains to intelligent agents to date and we briefly review the literature on the underlying queuing theory as it is applied in this paper. The rest of this paper is organized as follows. In Section 3, we use queuing theory to model agents when no manager is present and evaluate the expected revenue for the agent and the proportion of clients served. Because the agent must allocate its time between seeking and servicing clients, the server must decide whether or not to accept clients based upon expected revenue and service time. Lippman and Ross (1971) provide a decision rule (Section 3.1) for accepting clients for a single server in a market of heterogeneous clients. In Section 3.2, we extend their rule for a heterogeneous two-server system and note that the actions of each server will impact the expected revenue for both servers in the marketplace. Moreover, the decisions of each server impact the number of clients of each type served. We completely solve the balance equations for a two server, two client-type system and present a numerical example in Section 3.3 to illustrate the impact of these decisions. Then, in Section 4, we discuss the social impact and efficiency of the agents' decisions and we introduce a manager to demonstrate that both the clients and the agents are better off when a manager performs the roles of seeking and assigning clients. In Section 5, we offer conclusions and areas for further research.

2. Literature Review

We begin by briefly reviewing the literature on market-based control systems because it is the precursor to the emerging research on intelligent agents in electronic commerce. Then we present efforts to model a few specific aspects of intelligent agents in these markets (auctions) and note that, to date, no one has focused on how work-seeking agents should allocate time between searching for and performing tasks. Moreover, decision rules for what tasks intelligent agents should accept or reject have not been considered. Therefore, we review admission and rejection policy in queuing models in other domains. In the electronic commerce domain, there has been limited work on agents as intermediaries or brokers and we cite those examples. However, no one has focused on organizing intelligent agents in electronic commerce into groups headed by a manager agent as we are proposing in this paper.

2.1 Agents in Market-Based Control Systems and Auctions

Miller and Drexel (1988) provide a framework for applying market-based mechanisms to organize computational resources in large computer systems. The rationale for applying market mechanisms to complex resource allocation problems is that open, or argoric, markets have been allocating resources successfully for millenniums. Their work details many practical comparisons between social markets and their proposed computational organization. In addition to enumerating a number of potential roles for agents in a computational market, they emphasize the need for appropriate rules to coordinate desirable results. Clearwater(1996) provides further support for using market-based control as a technique for allocating resources in complex organizations. He also offers several practical examples of market-based control. Mullen and Wellman (1996) discuss a number of issues in the design of agents in computational markets.

Their focus is on theoretical aspects of market mechanisms and the application of economic principles to these new mechanisms.

Research has been done on modeling decision rules for intelligent agents that participate in auctions. For example, Parkes (1999) weighs the benefit of attaining more information with regard to the actual worth of a good versus the cost of determining that information for intelligent agents under various auction designs. Sandholm (1993) models autonomous agents that bid for delivery tasks performed by vehicles. Based upon their locally defined utility function and constraints, agents can calculate marginal costs and determine optimal bidding strategies in competitive and cooperative markets. Although Sandholm considers the time to calculate and recalculate optimal bids to be important in cost calculations, he assumes that deliveries are being accomplished simultaneously with auctions for future deliveries. In this paper, we address the problem of how an agent should allocate its time between seeking and accepting tasks, a similar problem that faces human's such as self-employed consultants. Moreover, a consultant must decide whether to accept a client or to continue seeking work in hopes of securing a more lucrative job.

2.2 Acceptance Policy for the Self-Employed

Lippman and Ross(1971) derived the acceptance decision rule (Section 3.2) for a scenario similar to the one faced by a self-employed consultant. This rule provides the basis for our analysis of a single agent's optimal behavior. Lippman and Ross model a semi-Markov decision process for a single streetwalker who must decide whether to accept customer offers based upon the ratio of offered reward to expected service time. In their model, preemption and backlogging are not permitted and customers who are rejected or find the streetwalker busy are lost. We apply their decision rule to a single intelligent agent and extend their result to a market where there are more than one server. In our two server model, in the absence of a manager, customers who find both servers busy are lost. When a manager is introduced, customers may be queued.

Other examples of acceptance policy analysis include, Carrizosa *et al* (1998) and Johansen (1994). Carrizosa *et al* model an emergency vehicle system as a M/G/c/c queue where customers that are not admitted to the system are serviced by a more costly backup system. This model differentiates customers by the distance vehicles must travel to service them. With minimizing costs as their objective, they derive an optimal acceptance policy of accepting an *i*-type customer when a server is available with probability x_i where the x_i 's are the control variables.

Johansen (1994) models a M/G/1 FIFO jobshop in which the owner must manage the backlog to: maximixe profits (when jobs come from an external source) or maximize welfare (when jobs are arriving from within a company). The backlog is important because jobs coming to the jobshop can evaluate their delivery time based upon the work in the system. Hence, a customer's net benefit is a function of the backlog, its own processing time, and the price. Customers enter the system if the net benefit is nonnegative.

2.3 Agents as Intermediaries

In general, electronic intermediaries may be thought of as performing roles that a broker, such as a real estate broker, does in traditional commerce. Resnick *et al* (1995) emphasize that electronic intermediaries are valuable because they can reduce buyer search costs, give privacy to buyers and sellers, provide insurance against bad behavior (similar to credit card companies), and match buyers and sellers. Similarly, Bailey and Bakos (1997) identified four intermediary roles: aggregation, trust, facilitation, and matching. They surveyed 13 electronic markets and asked them to assess these roles' importance and compare the traditional to the electronic intermediaries' capability to perform each role. They concluded that the degree to which electronic intermediaries can assume these traditional roles depends on the nature of the market.

Mullen and Wellman (1998) present a specific type of intelligent intermediary, an auction manager in electronic commerce. They describe the auction manager as an example of market middleware. In contrast to intermediaries and market middleware, we are interested in modeling the individual agent that must perform a number of tasks and exploring the possibility of a new type of agent, a *middle agent*, which can perform tasks that a manager in a traditional organization might perform. Some of these tasks may overlap with those of a broker such as providing insurance or matching buyers and sellers. But, unlike the brokers described above, our manager agent is responsible for a group of intelligent agents. In consultant's terms, our manager agent is like a principle director who finds work for the company and allocates it amongst a group of consultants. In this paper, we only model the task of assigning work to agents. However, we envision a middle manager that might perform many more tasks that are similar to those performed by a human middle manager such as supervision and quality control.

In summary, there has been a great deal of research on intelligent agents as participants in market-based control mechanisms and as intermediaries in electronic commerce. However, there has not been much work on how intelligent agents should be organized to perform in the electronic commerce domain. We propose that intelligent agents might reflect human organizations by organizing into groups under supervisory managers.

3. Modeling Agents

An electronic agent that must divide its time between servicing clients and seeking them in the marketplace may be modeled as a single server system that does not allow a queue. This modeling technique captures optimal performance characteristics for the agent because it allows for maximum throughput and, therefore, revenue for the agent. We emphasize that such an agent cannot simultaneously seek and serve clients. If agents had the ability to multi-task, we could consider a queuing system with preemption. As noted in the Introduction, it is important to remember that multi-tasking ability does not do away with the problem at hand; it simply permits preemption of tasks. Adding preemption will not change the basic insights of this paper and thus, we will keep the analysis herein simple by assuming no preemption.

If we denote p_0 as the proportion of time that the server is idle for a M/M/1/1 system, then for our purposes p_0 will be the proportion of time the agent is seeking work. Thus, an agent is busy $(1-p_0)$ proportion of the time at its exponential service rate, m, in a market where customers are arriving according to a Poisson process at rate l. We can solve a simple optimization problem to maximize throughput subject to the constraint that an agent's output cannot exceed input. Specifically,

Max: $(1-p_0)m$

st. $(1-p_0)m \leq p_0 l$

$0 \leq p_0 \leq 1$

Our decision variable, $p_0^* = \dfrac{m}{m+l}$, that maximizes throughput coincides with p_0 for the M/M/1/1 system.

When there is more than one agent and various client types we can still use a zero-capacity queue to model the system. However, in order to calculate expected revenues and determine acceptance policies agents need to know whether other agents will be accepting each of the other client types. This is because the acceptance of client types by each agent in the system will affect the arrival rates of those types of clients to all agents. Additionally, agent decisions to accept or reject clients will have an impact on the proportion of each client type served. Therefore, to examine the impact of acceptance policies for n client-types in a m agent-type market, we must solve 2^{m+n} sets of balance equations. Fortunately, we may make several assumptions that will reduce that number. For example, we may assume that every agent services at least one client type. We may further reduce the number of sets of balance equations by making use of the structure of the agent's decision rule for accepting clients. This will be detailed in Section 3.2.

3.1 Decision Rule for Accepting Clients

Lippman and Ross (1971) modeled a scenario in which a single server must decide whether to accept or reject heterogeneous clients based upon their respective reward to service time ratio. We can extend their result to the multiple heterogeneous server case. While the decision rule for each server to accept or reject clients remains the same, the actions of other servers in the system impact the expected revenue for servers in the system. Moreover, a server's decision to service clients that it is less efficient in servicing can have an adverse impact on other clients in the system because they find this server busy a larger proportion of the time. We demonstrate this in the sections that follow by examining the case of two agent types and two client types and by presenting a numerical example.

Lippman and Ross offer the following rule for accepting clients for the single server case. Accept client type i with reward R_i and service time s_i if $\frac{R_i}{s_i} \geq g$ where g is the optimal long run expected return per unit time. The intuition behind this rule is that the server should be willing to allocate time to a client type as long as the expected reward per unit time for that client type is greater than the optimal expected reward per unit time. For the single server case, the result from Lippman and Ross can be applied by ordering clients from highest to lowest reward to service ratio and determining when

$$\frac{R_N}{s_N} < E[\text{ return for servicing first } N-1 \text{ type clients}]$$

Then the server should accept the first N-1 type clients.

When other agents are in the system the server should still incorporate this rule for deciding whether or not it should accept various client types. However, the calculation of expected revenue in the face of competition is complicated by whether or not other agents are servicing clients, how efficient other agents are and how clients choose agents in the marketplace. As mentioned previously, there are 2^{m+n} possible combinations of policies. We may reduce the number of sets of balance equations that must be solved by ordering the clients for each agent and determining where the cutoff points should be based upon the actions of the other m-1 agents.

3.2 The Two Agent-Type, Two Client-Type Market

We present a two agent-type, two client-type market to analyze the impact of agent acceptance policies on the marketplace. The clients arrive to the marketplace at rates λ_j {j=1,2} according to a Poisson process. An agent-type i processes client type j at rate μ_{ij}. The processing times follow independent exponential distributions with mean, $s_{ij} = 1/\mu_{ij}$. We further assume that only one event, an arrival or service completion can occur at an instant in time. Thus, we have satisfied the properties of a birth and death process (Ross, 1996).

As discussed above, we model a zero capacity queue and assume that each agent accepts at least one type of client. Each agent i charges a different price for each type of service j, R_{ij}. We assume that each agent would prefer to service its own client type if it could only service one type of client based upon the reward to service ratio, $\mu_{ii}R_{ii} > \mu_{ij}R_{ij}$, $i \neq j$. Hence, we have four systems to examine of the possible combinations of agent acceptance policies.

We also make an important assumption regarding client behavior. We assume that each client type seeks its own agent type first before checking the other agent. We make price and processing rate assumptions that are consistent with this assumption. We assume that agent type i is faster and cheaper than agent type j for a job type i, $\mu_{ii} > \mu_{ij}$ and $R_{ii} < R_{ij}$, $i \neq j$. Thus, we can enforce a market discipline where client types seek their respective agent types first and only approach the other agent if they find their agent type busy. This market discipline simplifies our acceptance decision rules because each agent will be assured to have at least the revenue generated for accepting the higher reward to service ratio clients regardless of the other server's decision.

Under these market conditions, the decision rule for an agent to accept or reject clients is not affected by the other server's decision. However, its expected return for accepting both client types is affected by whether or not the other agent decides to accept both client types. Hence, if a third client type were introduced to the system, the decisions regarding the first two client types by each of the agents would be important to making decisions regarding the new client type.

The decision rule for each client is straightforward. An agent, i, should reject the j type client if

$$m_{ij}\mathrm{R}_{ij} < \frac{m_{ii}\ l_i\mathrm{R}_{ii}}{m_{ii} +\ l_i};\ \ i \neq j.$$

The four possible combinations are summarized in Table 1. Hereafter, we refer to these four cases by their numerals.

Table 1 Conditions for Each Possible Decision Combination

	Agent Action	Conditions for Agent One	Conditions for Agent Two
I.	Each agent serves own client type	$m_{12}\mathrm{R}_{12} < \frac{m_{11}\ l_1\mathrm{R}_{11}}{m_{11}+\ l_1}$	$m_{22}\mathrm{R}_{21} < \frac{m_{22}\ l_2\mathrm{R}_{22}}{m_{22}+\ l_2}$
II.	Agent two only serves type two clients, Agent one serves both	$m_{12}\mathrm{R}_{12} > \frac{m_{11}\ l_1\mathrm{R}_{11}}{m_{11}+\ l_1}$	$m_{22}\mathrm{R}_{21} < \frac{m_{22}\ l_2\mathrm{R}_{22}}{m_{22}+\ l_2}$
III.	Agent one only serves type one clients, Agent two serves both	$m_{12}\mathrm{R}_{12} < \frac{m_{11}\ l_1\mathrm{R}_{11}}{m_{11}+\ l_1}$	$m_{22}\mathrm{R}_{21} > \frac{m_{22}\ l_2\mathrm{R}_{22}}{m_{22}+\ l_2}$
IV.	Both agents accept both types of clients	$m_{12}\mathrm{R}_{12} > \frac{m_{11}\ l_1\mathrm{R}_{11}}{m_{11}+\ l_1}$	$m_{22}\mathrm{R}_{21} > \frac{m_{22}\ l_2\mathrm{R}_{22}}{m_{22}+\ l_2}$

To analyze the impact of each combination of decision policies we solved four sets of balance equations (Appendix A), one for each of the four cases. The solution is lengthy and does not lend itself to analytical interpretation; therefore we omit it. Instead, we present a numerical example consistent with the assumptions and market discipline outlined above to examine the four cases.

3.3 Numerical Example

To illustrate what happens when the marketplace is overwhelmed, we chose a numerical example that satisfies the condition, $l_1 > m_{11} + m_{21}$. Thus, the type one client could flood the market place by itself. However, it offers a relatively cheap reward, $\mathrm{R}_{i1}<\mathrm{R}_{i2}$. Meanwhile the type two client is relatively rare $\lambda_2 << \lambda_1$, but offers a large reward. We chose numbers for which pricing structures for all four decision combinations are feasible without violating these conditions and our previous assumptions.

$$l_1 = 30,\ l_2 = 1,\ m_{11} = 20,\ m_{12} = 2,\ m_{21} = 10,\ m_{22} = 4$$

3.3.1 Expected Revenue

We summarize our results, in terms of expected revenue per unit time for each agent type for the four cases, in Table 2 and discuss the effect agent decisions have on themselves and each other. How the agents' acceptance policy impacts the client types will be discussed in the next section of this paper.

Table 2 Expected Revenue per Unit Time

	Policy	**Agent One**	**Agent Two**
I.	Each agent serves own client type	12 R_{11}	.8 R_{22}
II.	Agent two only serves type two clients, Agent one serves both	11.158 R_{11} + .070138 R_{12}	.8 R_{22}
III.	Agent one only serves type one clients, Agent two serves both	12 R_{11}	4.7950 R_{21} + .320501 R_{22}
IV.	Both agents accept both types of clients	9.6973 R_{11} + .19189 R_{12}	5.0148 R_{21} + .29852R_{22}

We note that based on the market discipline we have defined when one agent type accepts both types of clients while the other only accepts its own type, the latter agent's revenue is unaffected. However, if an agent decides to service both types of clients, the expected revenue for the agent depends on the other agent's decision. Clearly, had we not assumed that clients seek their own server type first, our decision rules for this simple two client-type, two server-type system would be more complex. This demonstrates how decisions for agents can be complicated as the need to consider other agents and more client types increases.

It is also interesting to observe how an agent's revenue is impacted by the other agent's acceptance policy. For example, comparing Case II versus Case IV, we see that agent one's revenue for type one clients decreases when agent two decides to accept client type one, as expected. However, we see a simultaneous increase in the revenue generated for agent one by client type two. This is because agent two is accepting and servicing client type one inefficiently. Thus, it is unavailable for the type two client more often in Case IV than in Case II, thereby increasing the type one agent's revenue from type two clients.

4. Introducing the Manager Agent

In this section, we make some observations regarding the social impacts and efficiency of the current system and discuss the potential role of a central authority. Then we discuss how the introduction of a manager agent to this market would increase revenues for the agents combined with or without allowing a backlog of work to form. We close this section with a brief discussion of a number of roles in addition to seeking and assigning work that a manager agent could perform.

4.1 Social Impact and Efficiency

The social impact may be measured in terms of what proportion of the customers are served. Efficiency refers to whether or not customers are served by the most efficient server. Table 3 summarizes the proportion of client-types served by agent-type for each of our four cases. Numbers in parentheses are the total proportion served.

Perhaps the most interesting part of our analysis is the social impact of the self-interested agents' decisions. The consequences for type two clients are the most dramatic. If the type two agent only services type two clients, the type two clients realize an 80-87% acceptance rate. However, if it is cost effective for the type two agent to accept type one clients, over half the type two clients are lost and as many as 68% of the type two clients may be turned away. The worst case for clients type two is the best case for clients type one and vice versa. These two cases also represent the most and least number of clients serviced in the aggregate.

Table 3 Proportions of Each Client Type Served and Total Number Served per Unit Time

	Policy	**Proportion of type one clients served by agent type:** one/two (total)	**Proportion of type two clients served by agent type:** one/two (total)	**Total number served per unit time**
I.	Each agent serves own client type	.4 / 0 (.4)	0 / .8 (.8)	12.8
II.	Agent two only serves type two clients, Agent one serves both	.37195 / 0 (.37195)	.07014 / .8 (.87014)	12.03
III.	Agent one only serves type one clients, Agent two serves both	.4 / .15983 (.55983)	0 / .32050 (.32050)	17.12
IV.	Both agents accept both types of clients	.32324 / .16716 (.49040)	.19189 / .29851 (.49040)	15.20

In addition to the changes in proportion of clients served, recall that the clients that are served by the alternate agent type are serviced more slowly and at a higher price. In Case IV, over 50% of each client type served is serviced at the least efficient rate. Depending upon the goals of society, a central authority might wish to impose a particular decision combination for our agents. For example, the central authority might wish to have an efficient marketplace. To enforce a decision combination, as in Case I, the central authority might place restrictions on the agents through licensing or other forms of certification. Alternatively, society might wish to impose "fairness" criteria on the marketplace in favor of Case IV. Then the central authority could artificially change revenues through taxation.

4.2 Manager Agent

Similar to a central authority, manager agents could enforce a decision combination that the individual agents would not arrive at on their own. In the case of the manager agent, its objective would be to maximize revenues for the agents combined. As an example, consider our numerical example and revenues consistent with our assumptions of $R_{11} = 10$, $R_{21}=19$, and $R_{12}=R_{22}=200$. The self-interested agents would select Case II resulting in revenues of 125.608 and 160.000 per unit time for agents one and two respectively. However, a manager agent could realize a total of 290.336 per unit time for the agents combined if it enforced an agreement between the agents to accept both types of clients (Case IV). A manager agent would be necessary to enforce such collusion because the type two agent by itself is better off if it only accepts type two clients. Thus, without changing the fact that agents do not allow queuing, a manager agent can increase revenues for the agents combined.

If a manager agent performed the search role for agents, this would be akin to allowing a queue. Agent "idle time" represents the search time for our electronic agents. It also represents time that the agent could be working if it had a manager agent that found work for it. We note that the idle time is significant under all four cases (Table 4), particularly in light of the fact that type one clients could flood the market if queuing were allowed. If a manager agent were introduced in the market to manage queues for these agents, it could allocate clients to reduce idle time for the agents and increase overall market efficiency. This is true, in general, and it is clearly illustrated by our numerical example.

Table 4 Expected Proportion of Time Each Agent is "Idle"

	Policy	Agent One	Agent Two
I.	Each agent serves own client type	.4	.8
II.	Agent two only serves type two clients, Agent one serves both	.37122	.8
III.	Agent one only serves type one clients, Agent two serves both	.4	.32050
IV.	Both agents accept both types of clients	.32324	.29852

Recall that each agent's revenue to service ratio is higher for its own client type. In addition, type one clients could flood the market. A manager agent could maximize revenues for the agents and limit market inefficiencies in the following way. It could ensure that agent type one is busy 100% of the time with client type one. As for the second agent, the manager would want to ensure that agent two never miss an opportunity to serve a type two client. It could accomplish this objective by giving type two clients a priority in the queue. Thus, for our numerical example it would be possible to ensure that 100% of the type two clients are served and 91.67% of the type one clients are served.

In our simple model, we only modeled the seeking and assigning work roles of a manager agent. However, there are a number of other manager tasks a manager agent could perform in an electronic market. For example, allowing a backlog introduces a new kind of inefficiency. In the absence of a manager agent, the clients that were served experienced no waiting time. In the arrangement described above, type one clients could experience infinite waiting times if the manager does not introduce an acceptance policy to limit waiting times. Determining an acceptance policy that optimizes revenue is one role a manager could perform. One strategy would be to create a finite queue. Another might be to introduce pricing schemes that decrease the potential arrival rates. This is an important area for further research that we do not detail here.

Related to managing customer waiting times is ensuring quality control for customers. This is one of many supervisory type tasks a manager agent could perform. Other supervisory tasks a manager agent could perform include: determining work and maintenance schedules, hiring more agents, and identifying needs for upgrades in customer service. Two additional important roles a manager agent could fill are bidder and entrepreneur. In the next few paragraphs we discuss each of these latter roles in turn.

As mentioned in the first section of this paper, intelligent agents can participate in auctions. As the electronic market place expands there will be an increase in the number of intelligent agents in the marketplace. This will have several implications. First, the auctions will become more competitive and will take longer to execute. For the single agent, this represents an increase in the time spent away from executing orders to gain a particular job. At the same time, with so many intelligent agents searching the marketplace the window of opportunity to place bids may be shorter than the amount of time necessary for the intelligent agent to calculate an efficient bid. A manager agent could perform this task, bidding, for the single agent. Moreover, the manager agent could position itself to observe the marketplace and place more efficient bids. For example, the manager agent could determine that the market is flooded and respond by raising prices.

As the amount of business conducted via intelligent agents increases, the single agent may find itself overwhelmed by opportunities to consider. As an entrepreneur, a manager agent could observe the marketplace and determine when such opportunities are worth pursuing. Furthermore, the manager agent could observe unmet demands for services and create new opportunities for the agents.

5. Conclusions and Further Research.

The purpose of this paper was to introduce the efficacy of a manager agent in the electronic market place. A manager agent can increase revenues for the intelligent agents by assuming the role of seeking clients and assigning them to the most efficient agent. The clients, in a congested market, benefit because their access to service is increased. However, those clients that are served experience waiting times that they did not experience in the absence of a manager. We believe that determining how the manager will manage the queue is an area for further research.

Areas for further research regarding intelligent agents in electronic commerce may be classified as structural issues or as design issues. The presence or absence of a manager agent in the marketplace is a structural issue. An important area for further research is determining when a manager agent is cost effective for the agents. Another key question is when a manager agent or a central authority might be effective or desirable in improving societal welfare. As discussed earlier, this depends on society's goals. Our numerical example provides an obvious situation where a manager would be cost effective in seeking and assigning work. As electronic commerce grows, agents may need a manager agent to perform other tasks as well. We believe this paper is a first contribution toward identifying what some of those tasks might be.

The issues regarding how the manager performs these tasks are design issues. A question such as how the manager enforces queue discipline in our example is a design issue. Another design issue is what policies a central authority might impose to achieve a societal objective. In sum, as we research the potential future of electronic commerce, we must explore both what roles the agents in the market will assume and how they will do them.

References

Alsop, S. “Is there an Amazon.com for every industry?” *Fortune*, January 11, 1999, 159-160.

Bailey, J.P. and Y. Bakos, “An Exploratory Study of the Emerging Role of Electronic Intermediaries,” *International Journal of Electronic Commerce*, Volume 1, No. 3 (1997), 7-20.

Bakos, Y. “Towards Friction-Free Markets: The Emerging Role of Electronic Marketplaces on the Internet,” *Communications of the ACM*, 41, (August 1998), 35-42.

Bakos, J.Y., “Reducing Buyer Search Costs: Implications for Electronic Marketplaces,” *Management Science*, 43,(1997),

Carrizosa E., E.Conde, and M. Munoz Marquez, Admission Policies in Loss Queueing Models with Heterogeneous Arrivals. *Management Science*, 44, (1998),311-320.

Chavez, A. A. Moukas, and P. Maes, “Challenger: A Multi-agent System for Distributed Resource Allocation,” Agents ’97 Conference Proceedings, 1997.

Clearwater, S.H. Editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*, World Scientific, Singapore ,1996.

Clearwater, S.H. and B.A. Huberman, “Thermal Markets for Controlling Building Environments,” Energy Engineering, 91,No. 3 (1994), 26-56.

Duncan, E. “Technology and entertainment: A digital future,” *Economist*, November 21, 1998, S14-S16.

Fox, B. “Does the World Still Need Retailers?” *RTMagazine*, April 1999.

Gellman, R. “Disintermediation and the Internet,” *Government Infromation Quarterly*, 13, No. 1, 1-8.

Johansen, Soren Glud, Optimal Prices of an M/G/1 Jobshop, *Operations Research*, 42,(1994) 765-774.

Klienrock, L. *Queueing Systems, Volume 1: Theory*, John Wiley, New York, N.Y., 1975.

Lippman, S. A. *Applied Probability Models with Optimization Applications*, Holden Day, San Francisco, CA, 1970.

Lippman, S. A. and S. Ross, “The Streetwalker’s Dilemma: A Job Shop Model,” *SIAM J. Applied Math*, 20 (1971), 336-342.

Malone, T.W., J. Yates and R.I. Benjamin, “The Logic of Electronic Markets,” *Harvard Business Review*, May-June 1989, 166-170.

Miller, M.S. and K.E. Drexler, “Markets and Computation: Argoric Open Systems,” In B.A. Huberman (ed) *The Ecology of Computation,* Elsevier Science, North-Holland, 1988.

Mullen, T. and M.P. Wellman, “The Auction Manager: Market Middleware for Large Scale Electronic Commerce,” *Proceedings of the Third USENIX Workshop on Electronic Commerce*, Boston, MA, September 1998.

Mullen, T. and M.P. Wellman, “Some Issues in the Design of Market-Oriented Agents,” In M. Wooldridge, J. Mueller, and M. Tambe (eds.*), Intelligent Agents II: Agent Theories, Architectures, and Languages*, Springer-Verlag, 1996.

Organisation for Economic Co-operation and Development, *The Economic and Social Impacts of Electronic Commerce: Preliminary Findings and Research Agenda*, OECD, Paris, France, February 1999.

Parkes, D., “Optimal Auction Design for Agents with Hard Valuation Problems,” In *Proceedings of the 2nd Workshop on Agent Mediated Electronic Commerce* (AmEC-99). Stockholm.

Resnick, P. , R. Zeckhuaser, and C. Avery, “Roles for Electronic Brokers,” *in Toward a Competitive Telecommunication Industry: Selected Papers from the 1994 Telecommunications Policy Research Conference*, G.W. Brock, (Ed), Mahwah, NJ: Lawrence Erlbaum Associates, 1995, 289-304.

Ross, S. *Stochastic Processes*, 2nd Edition, John Wiley, New York, N.Y., 1996.

Sandholm, T. “An Implementation of the Contract Net Protocol Based on Marginal Cost Calculations;” In *Proceedings of the National Conference on Artificial Intelligence*, pages 256-262, Washington, D.C., 1993, AAAI.

Secretariat on Electronic Commerce, *The Emerging Digital Economy*, U.S. Department of Commerce, Washington, DC, April 1998.

Shapiro, C. and H.R. Varian , *Information Rules: A Strategic Guide to the Network Economy*, Havard Business School Press, 1999.

Swaminathan, J.M., S. F. Smith, and N. M. Sadeh, “Modeling Supply Chain Dynamics: A Multiagent Approach,” *Decision Sciences*, 29 (1998), 607-632.

Taylor, P. “Middle men deleted as word spreads,” *Financial Times*, October 27, 1998, 15.

Wilder, C. “Middlemen beware?” *Informationweek*, October 20, 1997, 94-98.

Appendix A

Balance Equations and Transition Diagrams

There are nine possible states for case IV. So we define the states for this system and use the appropriate subset for the other three systems.

State klmn k=1 agent one is servicing type one client, 0 otherwise

l=1 agent one is servicing type two client, 0 otherwise

m=1 agent two is servicing type one client, 0 otherwise

n=1 agent two is servicing type two client, 0 otherwise

If we let P_{klmn}= the probability of the market being in state klmn, we may solve the balance equations for the probabilities and derive the expected revenues and proportion of each client type served.

Case I: Each agent serves own client type.

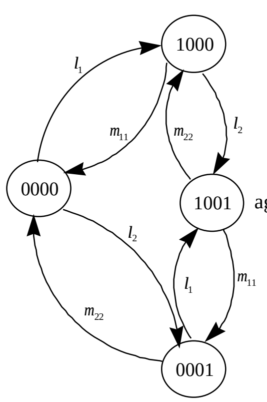

Case I: Each agent serves own client type.

$$\left(l_1 + l_2\right)\mathrm{P}_{0000} = m_{11}\mathrm{P}_{1000} + m_{22}\mathrm{P}_{0001}$$
$$\left(l_2 + m_{11}\right)\mathrm{P}_{1000} = l_1\mathrm{P}_{0000} + m_{22}\mathrm{P}_{1001}$$
$$\left(m_{11} + m_{22}\right)\mathrm{P}_{1001} = l_2\mathrm{P}_{1000} + l_1\mathrm{P}_{0001}$$
$$\sum_{\substack{\text{all klmn} \\ l\neq 1,\ m\neq 1}} \mathrm{P}_{\text{klmn}} = 1$$

The solution for the four possible states for Case I is

$$\mathrm{P}_{0000} = \tfrac{l_1\, l_2}{h};\quad \mathrm{P}_{1000} = \tfrac{m_{11}\, l_2}{h};\quad \mathrm{P}_{1001} = \tfrac{m_{11} m_{22}}{h};\quad \mathrm{P}_{0001} = \tfrac{m_{22}\, l_1}{h};\quad h = \left(m_{11} + l_1\right)\left(m_{22} + l_2\right)$$

The solutions for the remaining three cases are cumbersome and not conducive to analysis and so we omit them.

Case II: Agent two only serves type two clients, Agent one serves both.

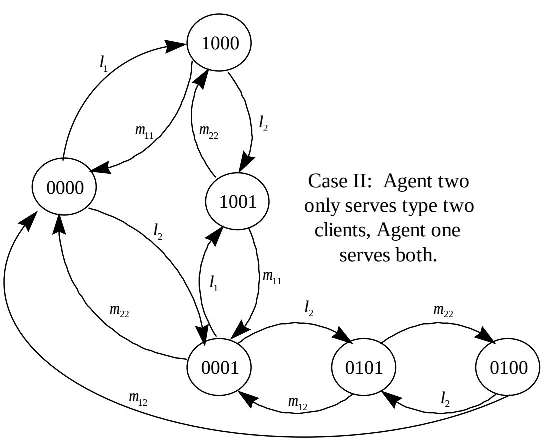

$$\left(l_1 + l_2\right)\mathrm{P}_{0000} = m_{11}\mathrm{P}_{1000} + m_{22}\mathrm{P}_{0001} + m_{12}\mathrm{P}_{0100}$$
$$\left(l_2 + m_{11}\right)\mathrm{P}_{1000} = l_1\mathrm{P}_{0000} + m_{22}\mathrm{P}_{1001}$$
$$\left(m_{11} + m_{22}\right)\mathrm{P}_{1001} = l_2\mathrm{P}_{1000} + l_1\mathrm{P}_{0001}$$
$$\left(m_{12} + m_{22}\right)\mathrm{P}_{0101} = l_2\mathrm{P}_{0001} + l_2\mathrm{P}_{0100}$$
$$\left(m_{12} + l_2\right)\mathrm{P}_{0100} = m_{22}\mathrm{P}_{0101}$$
$$\sum_{\substack{\text{all klmn} \\ m\neq 1}} \mathrm{P}_{\text{klmn}} = 1$$

Case III: Agent one only serves type one clients, Agent two serves both.

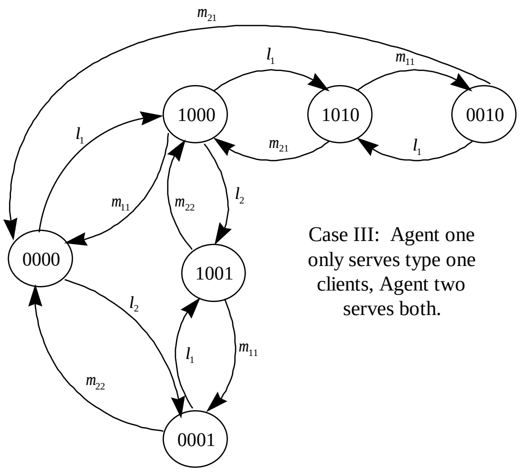

$$\left(l_1 + l_2 \right) \mathrm{P}_{0000} = m_{11} \mathrm{P}_{1000} + m_{22} \mathrm{P}_{0001} + m_{21} \mathrm{P}_{0010}$$

$$\left(m_{11} + m_{22} \right) \mathrm{P}_{1001} = l_2 \mathrm{P}_{1000} + l_1 \mathrm{P}_{0001}$$

$$\left(l_1 + m_{22} \right) \mathrm{P}_{0001} = l_2 \mathrm{P}_{0000} + m_{11} \mathrm{P}_{1001}$$

$$\left(m_{11} + m_{21} \right) \mathrm{P}_{1010} = l_1 \mathrm{P}_{1000} + l_1 \mathrm{P}_{0010}$$

$$\left(l_1 + m_{21} \right) \mathrm{P}_{0010} = m_{11} \mathrm{P}_{1010}$$

$$\sum_{\substack{\text{all klmn} \\ \text{l} \neq 1}} \mathrm{P}_{\text{klmn}} = 1$$

Case IV: Both agents accept both types of clients.

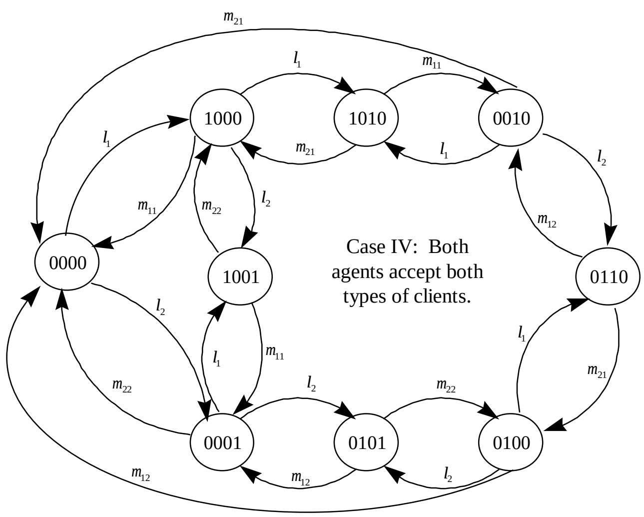

$$\left(l_1 + l_2 \right) P_{0000} = m_{12} P_{0100} + m_{11} P_{1000} + m_{21} P_{0010} + m_{22} P_{0001}$$

$$\left(l_1 + l_2 + m_{11} \right) P_{1000} = l_1 P_{0000} + m_{21} P_{1010} + m_{22} P_{1001}$$

$$\left(m_{11} + m_{22} \right) P_{1001} = l_1 P_{0001} + l_2 P_{1000}$$

$$\left(l_1 + l_2 + m_{22} \right) P_{0001} = l_2 P_{0000} + m_{11} P_{1001} + m_{12} P_{0101}$$

$$\left(m_{11} + m_{21} \right) P_{1010} = l_1 P_{0010} + l_1 P_{1000}$$

$$\left(m_{12} + m_{22} \right) P_{0101} = l_2 P_{0001} + l_2 P_{0100}$$

$$\left(l_1 + l_2 + m_{21} \right) P_{0010} = m_{11} P_{1010} + m_{12} P_{0110}$$

$$\left(l_1 + l_2 + m_{12} \right) P_{0100} = m_{21} P_{0110} + m_{22} P_{0101}$$

$$\left(m_{12} + m_{21} \right) P_{0110} = l_2 P_{0010} + l_1 P_{0100}$$

$$\sum_{\text{all klmn}} P_{\text{klmn}} = 1$$