

Congregation Formation in Information Economies

Christopher H. Brooks and Edmund H. Durfee

Artificial Intelligence Laboratory
University of Michigan
Ann Arbor, MI 48109
{chbrooks, durfee}@umich.edu

Abstract

In a large-scale multiagent system, agents that need to interact with other agents are faced with a combinatorially explosive number of potential interactions. One way for agents to deal with this complexity is to form *congregations*. As in human society, congregations provide a common meeting-place for agents with compatible needs or preferences. We discuss two approaches to inducing congregations to form in a general multiagent system: external mechanisms and internal learning. We then examine a particular multiagent system, an information economy, in more detail, and discuss the relationship between bundling of information goods and the formation of congregations. By viewing the problem of determining the optimal bundling strategy as a problem of congregation formation, new techniques, such as ontological information about consumer preferences and goods, may be brought to bear. We also present some preliminary experimental results regarding both the conditions which lead to optimal congregation formation and the ability of a producer to learn some simple preferences of a congregation.

Introduction

In a multiagent system, most agents typically will not interact with every other agent. Rather, they will interact with a subset of the agents in the system. Ideally, this subset will be composed of agents with complementary needs, goals or preferences. That is, agents will tend to group together with agents that share some important features and avoid interacting with agents which do not share these features.

This grouping together is referred to as a *congregation*. Specifically, a congregation is a group of agents that has come together for some (ideally) mutually beneficial purpose. This may be the exchange of goods, services, or information, the accomplishing of tasks, or an aggregation in order to accomplish goals which could not be met separately. Congregations are typically denoted in some way that is meaningful to the agents in a system. This allows an agent to reduce the amount

of time spent searching for suitable ‘partners.’ In some cases, an agent may belong to more than one congregation, depending on its needs and the particular characteristics of the other agents in the society.

The forming of congregations is a common phenomenon in human society: people seem to be quite adept at discovering the proper people to deal with, even though they may be complete strangers. The congregations that exist in a society (clubs, markets, churches, etc.) make the problem of finding people with complementary interests or capabilities much easier for humans.

Our question becomes: how can we emulate this behavior within a computational framework? In particular, how can we place a useful sort of structure on the space of agent congregations which makes it easy for agents to find suitable ‘partners?’ In this case, we are interested in (ideally) finding configurations that are acceptable both to the agents themselves (i.e. have high local utility) and desirable from a global standpoint. (For example, efficient, or globally optimal, or maximal in some other sense, depending on the domain of interest.)

An economy is an example of a multiagent system in which congregations naturally occur, typically in the form of markets. Consumers with a particular set of needs will gather together in common locations to buy from like-minded producers. A farmer’s market is an example of this: a number of sellers of produce gather together to attract a larger number of consumers. Likewise, the consumers’ problem of finding an appropriate seller is simplified: they know that they need only go to the farmer’s market in order to find the particular fruit they are seeking.

There is a great deal of domain knowledge that consumers and producers are using in this decision to go to a farmer’s market to buy or sell produce, rather than a different congregation, such as a grocery store, a county fair or a restaurant. This seems to be common in human culture; people use a great deal of knowledge about the structure and expected characteristics of a group to simplify their decisions about which producers and consumers to deal with.

In an information economy, the producers and con-

This work was supported, in part, by the National Science Foundation under grant IIS-9872057 and by a grant from the IBM Corporation.

sumers may be software agents. In such an environment, they will need to make similar decisions about which other agents to buy and sell from. Consequently, we would like them to be able to quickly and easily decide which congregations are beneficial for them to join and which are not. Just as human congregations are described with identifiers that ease this search problem, we would like to describe computational congregations in a way that makes it easy for agents to identify congregations that they are particularly well-suited for.

In this research, we are interested in a particular type of congregation formation: the grouping together of consumers and producers of information goods in such a way as to produce a desirable global state. In an economy containing producers and consumers with heterogeneous preferences, the producers may end up segmenting the market into specialized niches in which each producer interacts with a subsets of consumers having particular interests, rather than all producers competing amongst themselves for the business of every consumer. In economies with multiple producers, this specialization into congregations will reduce the competitive burden on the producers while making it easier for a consumer to purchase goods that more closely match its preferences. In a monopoly, identifying the appropriate label for a congregation of consumers can be seen as a method for extracting the maximum profit, as is discussed below.

In particular, we are interested in relating the bundling of information goods to congregation formation. By choosing a particular subset of information goods to bundle together, a producer can change the makeup of the consumers which purchase its goods. How can producers choose an optimal bundling strategy so as to maximize either their own profit or to maximize some global measure, such as total social welfare?

Most of these questions remain open. In this paper, we present an approach to solving this problem. We begin by describing some related projects and their results. We then describe the general problem of congregation formation and discuss some strategies for helping agents find optimal configurations. Next, we describe the specific scenario of producers constructing congregations through the choice of bundles to be sold and present some preliminary results for a single-producer market. Finally, we conclude with some observations regarding future directions for research.

Related Work

Many of the ideas in the Congregating Agents project have grown out of the Service Market Society (SMS) (Durfee *et al.* 1998). This project was developed as an agent infrastructure for the University of Michigan Digital Library (UMDL) project. It featured a dynamically changing society of agents providing a variety of services, along with protocols to facilitate their communication and cooperation. Agent services were sold in a number of auctions. When an agent needed a particular service, it would consult a Service Classifier

Agent, which would direct it to the appropriate auction. As auctions grew and new services appeared, new auctions were formed and old auctions disappeared. Our work generalizes this, abstracting away from the details of auction mechanisms and treating any grouping of agents (including an auction) as a congregation. Once agents have congregated, we assume that they can use some mechanism, such as an auction, to exchange goods and services. This allows us to focus on the questions of under what conditions should new congregations form, how congregations are described to prospective members, what specific congregations should be created in a particular setting, and how the choice of congregations affects overall system performance.

Kephart *et al.* examine the dynamics of an information economy in (Kephart, Hanson, & Sairamesh 1998; Kephart *et al.* 1998). They consider a model in which brokers must choose a price and category of information goods to provide to consumers. They show that, in their model, when brokers react myopically to the present situation, the system will oscillate through an endless series of price wars. In (Tesauro & Kephart 1998), they provide some hope that if brokers employ some sort of lookahead, this can reduce system oscillations and lead to a stable configuration. While many of the details of this model are different from our own, it serves to inform us that agents which are capable of reasoning about and learning from their performance are one potential mechanism for finding a stable (and hopefully optimal) set of congregations.

Hogg and Huberman (Hogg & Huberman 1991) take a different approach to the problem of controlling the global behavior of a large-scale multiagent system. They show that properly applied external rewards and penalties can be used to stabilize an otherwise chaotic system. We take this as encouragement that external pressures can indeed be applied to an information economy to help agents discover congregations which have globally desirable characteristics.

Congregating Agents

As noted above, the focus of our work is on congregation formation. In its most general form, we want to abstract away from the details of buying and selling (such as auction mechanisms, the costs/benefits of deception, and so on) and focus on the question of how and whether agents which need to interact with each other can find each other amidst a large number of other, less interesting agents. We are currently considering two families of approaches to this problem: First, the imposition of external, system-designed mechanisms designed to encourage or discourage particular behaviors, and second, the use of learning and explicit modeling of agents and their preferences to construct the appropriate set of congregations.

Transaction Fees

External mechanisms for controlling agent behavior have the advantage that no *a priori* assumptions about

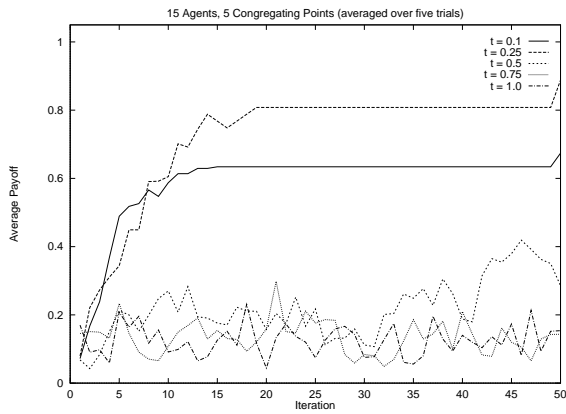


Figure 1: Results of Experiment 1 (t : transaction fee)

an agent’s internal reasoning mechanisms need to be made. As long as the agents behave rationally, rewards and punishments can be constructed so as to encourage and discourage certain types of behavior. Of course, the difficulty of this construction may depend upon a number of factors, including the complexity of the domain and the number of agents.

One simple external mechanism for controlling agents’ search for the correct congregation is to impose some sort of tax or transaction fee. On every iteration, this fee is subtracted from an agent’s payoff. This serves to make suboptimal congregations less appealing to an agent, since the cost of joining the congregation may be greater than any payoff received. This method of control is extremely general; it does not require agents to share any sort of knowledge or be able to reason about their knowledge. It only assumes that they will avoid states in which they are penalized more than they receive as a reward.

We have conducted some preliminary experiments which demonstrate the usefulness of a transaction fee in inducing agents to move to the correct congregation or, in more complex scenarios, at least avoid moving to a particularly bad congregation. In these experiments, a set of n affinity groups of agents were generated. Each agent prefers to congregate with other agents of its affinity group. A set of n congregating points were also generated, and agents were dispersed randomly throughout these groups. On each iteration, an agent determined its payoff from being in a congregation and subtracted the transaction fee. If the result was non-negative, it would stay in place. Otherwise, it would randomly move to a different congregating point.

In the first experiment, 15 agents were divided into five affinity groups, with five congregating points. Payoffs were an average of the pairwise interaction with every other agent in a congregation; an agent receives a 1 for interacting with members of its affinity group and a 0 for interacting with a member of a different affinity group. Results for different transaction fees, averaged over five trials, are shown in figure 1.




Figure 2: Results of Experiment 2 (t : transaction fee)

In this case, there is a clear value for the transaction fee, slightly above $\frac{1}{n}$, at which the problem of finding the correct set of congregations becomes easy. At $t = 0.25$, the agents converged to the globally optimal configuration four out of five times, finding a slightly suboptimal configuration on one iteration. At $t = 0.25$, most configurations containing two or more agents in the same affinity group will result in positive payoffs for these agents. They will therefore “hold still”, giving the other members of their affinity group a chance to find them. Above this value, the system thrashes about, as agents chase each other through different configurations of congregations. Below this value, the system quickly stabilized, but often settled into a local (and non-optimal) maximum.

The second experiment replicated the setup of the first experiment on a larger scale, using 30 agents and 10 affinity groups. The results are depicted in figure 2.

Once again, there is a transaction fee near $\frac{1}{n}$ at which the system quickly converges. Unlike the previous experiment, it does not reach the global optimum. In fact, the experimental data seem to suggest that in this case, a single transaction fee is not sufficient to force the system into a global optimum. The system does, however, quickly move away from states which are deemed to be globally undesirable.

In the third experiment, the nature of the payoff function was changed. Rather than being strictly Boolean, the differences between the affinity groups was fractional. For example, an agent of group A would receive a payoff of 0.75 from interacting with an agent of group B, 0.5 for interacting with an agent of group C, and 0.25 from interacting with group D. These fractional payoffs can be interpreted as partial similarity (i.e., the two affinity groups are similar but not identical) or the production of goods that either satisfy a fraction of an agent’s needs, or that only satisfy an agent’s needs a portion of the time. Once again, 30 agents and 10 affinity groups were generated. The results are shown in figure 3.

In this case, the transaction fee has very little ef-




Figure 3: Results of Experiment 3 (t : transaction fee)




Figure 4: Results of Experiment 4

fect. This is no doubt related to the increase in the number of (distinct) congregations and the varying effects that agents from different affinity groups have on each other. These fractional payoffs greatly increase the space of congregations to be searched, while making non-optimal configurations more difficult to distinguish from optimal ones.

The fourth experiment expands upon the third. Fractional payoffs are retained. However, the transaction fee is gradually increased over time, starting at 0.01 and increasing at increments of 0.1 every 10 iterations. Figure 4 presents a graph of the average payoff for each market over time. Note that in this graph, the performance of individual congregations is presented, whereas in previous graphs the performance of all congregations for a given transaction fee was averaged.

Interestingly, in this experiment, the approach of gradually increasing the transaction fee seems to work up to a certain point, approximately $t = 0.7$. At this point, the system becomes chaotic and oscillates wildly. It seems that initial increments serve to move the system between local maxima; the transaction fee induces a small subset of the agents to move, thereby allowing the system to retain some of its stability. Once the

transaction fee is high enough, nearly all of the agents try to move, and no congregational structure is retained between iterations. If we average the performance of these markets when $t > 0.75$, we get a graph very much like that of experiment 3.

These experiments serve to show that for small, well-structured problems, a simple external mechanism such as a transaction fee can induce a set of agents to find an optimal set of congregations. However, the quality of the results declines as the complexity of the system increases. In these cases, external mechanisms may not be able to push the system all the way to a global optimum. This does not mean that they are useless, however. They should still be able to prevent the system from reaching configurations that are particularly undesirable, thereby making the problem easier for the agents involved. It seems most likely that, in a complex multiagent system, external mechanisms will need to be combined with methods that allow agents to learn the proper set of congregations to form.

Using Ontological Information

In the previous experiments, agent preferences were particularly simple; agents merely wanted to be with other agents of a similar “type.” In an actual multiagent system, an agent’s preferences are not merely random; rather, they are a function of the world it inhabits. This world, and these preferences, can be described in terms of an ontology (Weinstein & Birmingham 1998). For example, in the SMS, information goods and abilities were described in ontological terms, using description logic. This additional information regarding the structure and relationship between agents’ preferences should be usable, either by agents or system designers, in determining who a particular agent should congregate with. While our current work does not use explicit ontological information, we are moving in that direction by attempting to learn a few key parameters that best characterize a congregation.

The assumption that was made in the SMS was that agents would begin with a core ontology describing the essential facts about the world. Differing experiences, capabilities, or interactions would lead to different concepts being developed and ontologies growing apart. For example, most scientific disciplines develop terminology to refer to specific concepts. It may even be the case that different fields use different terms to refer to the same concept.

Two problems then arise with respect to congregations: First, can differentiated ontologies be reconciled to a common set of terms so as to allow members different congregations to interact? This question is the focus of (Weinstein & Birmingham 1998) and will not be treated here, except to comment that methods exist for mapping differentiated ontologies onto a common set of terms. Second, what term should be used to describe a congregation so as to best attract the ‘correct’ agents? In an economic setting, we might rephrase this question as: how should a producer choose a segment of a mar-

ket so to attract the ‘right’ consumers for its goods? In the SMS, markets often contained multiple producers, and the term which described their contents was the most general term describing all the producers that were present. There is no *a priori* reason for assuming this is the optimal choice. In some circumstances, a term that imperfectly describes some agents in a congregation while describing others more accurately and specifically may serve to attract a more optimal set of agents. In general, it is hoped that describing a congregation using some sort of ontological information will prove to be more useful than no information at all.

Bundling in an Information Economy

We now depart from the general discussion of congregation formation to deal with the formation of congregations in a particular domain, namely that of an information economy. In particular, we wish to examine the effects of bundling information goods on the formation of particular congregations of consumers. In this work, we will be interested in a slightly more limited question: whether a producer can learn the correct way to characterize a congregation of consumers. This can be viewed as a simplified version of the preference-learning problem discussed above. If a producer can learn the valuations placed on articles by a set of consumers, it can then determine the optimal pricing strategy.

Bundling within an information economy has received a great deal of attention in the economic literature (Bakos & Brynjolfsson 1998; Chuang & Sirbu 1998; Fay & MacKie-Mason 1998). The primary way in which bundling of information goods differs from the bundling of traditional goods (where bundling is typically not preferred) is that information goods have a negligible marginal cost to reproduce. This can lead to scenarios in which a producer may package together large bundles of information (such as a journal) even though each consumer may only be interested in a small subset of the information. Since the information goods can be reproduced without cost, the bundle as a whole can be sold cheaply, appealing to a wider variety of consumers and earning the producer a larger profit.

Most economic analysis has focused on situations with a monopolistic producer, and seems to indicate that bundling can be an effective strategy. In our research, we wish to focus on the question of how multiple producers can choose the correct bundle size and content. By doing this, they are creating different congregations of consumers who wish to buy from them.

Congregation Formation Through Bundling Strategies

Consumers typically vary in the types of information goods that they are interested in, the quantity of goods desired, and the values they place on these goods. Since these goods have some semantic content, we should be able to find an ontological description characterizing the sorts of goods a consumer prefers. For example,

given an article about Michael Jordan’s retirement, one consumer might place a high valuation on any article about the Chicago Bulls, and so value this article highly, while another consumer might place a lower value on all articles about the NBA, thus valuing the article slightly less.

This brings up a related point. In traditional economic models, goods are frequently treated as complements or substitutes for each other independent of the valuations of the consumers present. For example, one good might be said to be 70% substitutable for another good. Additionally, all information about this good is compressed into a single numeric value. This may pose a problem for a producer; when deciding whether to bundle information goods together, it must consider the preferences of the consumers. In the previous example, the second consumer would likely be willing to substitute an article about the Los Angeles Lakers for the Michael Jordan article, while the first would not.

Both consumer preferences and the characteristics of information goods can be expressed in terms of an ontological structure characterizing those features that are deemed important. Some possible features include: size, cost, content, author, target audience, accompanying graphics, and so on. Since an economic congregation is composed of one or more producers and the consumers who purchase their goods, the problem of finding the optimal set of bundles for producers to provide can be cast as a search for the optimal set of market niches or congregations. Specifically, given a set of information goods, a producer must choose a subset to bundle. This subset will be describable, and this description will attract a particular group of consumers, depending on both their preferences and the bundles being offered by other producers. Which bundle, and consequently which description, will attract the ‘best’ consumers and produce the most profit? And, on a global scale, can all producers find a set of bundles such that all consumers are able to purchase information goods that satisfy their preferences and all producers are able to profit?

This seems to occur in an *ad hoc* way in human economies. Magazines, cable channels, and websites all choose a particular content, format and appearance with the goal of serving a particular target audience. We wish to achieve these results in a computational economy.

Preliminary Experiments

As a first experiment, we have asked a very simple question: given a single congregation, can a producer find the bundle that best characterizes its preferences? In fact, we have simplified the problem even further; there are no substitutes or complements, and consumers vary only in the number of goods they value and the value they place on each good. These are considered to be the consumers two preferences of interest. The producer’s problem then becomes one of finding the appropriate bundle size and price for that bundle.

There are two reasons for these simplifications. The first is to connect our work to the analytical results produced in the economic literature, in particular the results in (Chuang & Sirbu 1998; Fay & MacKie-Mason 1998). Most economic works on bundling tend to treat the problem as a static game with a monopolist producer. This lets us at least compare our empirical results to those of similar analytic models.

The second is more pragmatic: in order to construct a complex simulation of an economy, one must proceed in steps. Solutions that are appropriate for one set of assumptions may not hold when these assumptions are changed. In the transaction fee experiments described above, some simple changes to the model introduced very different types of behavior. Therefore, the best way to proceed seems to be to begin simply and understand the relevant phenomena regarding a simple system and then add complexity incrementally, seeing where predictions fail, and adjusting accordingly.

Model description

This model is quite similar to the one described in (Chuang & Sirbu 1998; Fay & MacKie-Mason 1998). It consists of a set of N information goods, one producer, and c consumers. Each consumer c_i places a valuation on the n th most preferred information good as follows:

$$V(n)_i = \begin{cases} w_{0_i} \left(1 - \frac{n}{k_i N}\right) & \text{if } n < k_i N \\ 0 & \text{otherwise} \end{cases}$$

Consumers vary in their choice of w_0 (the value placed on the most preferred article) and k_i (the number of articles that have a positive valuation). The producer’s problem is to choose a bundle size and price that optimizes its profit. For large numbers of consumers, this is equivalent to finding a bundle size and price that maximizes profits over the mean consumer, or, in other words, finding an efficient characterization of the preferences of the congregation of consumers, where preferences consist of article valuations and bundle sizes.

Experimental Results

In our initial experiments, a producer providing a set of 1000 articles to a potential congregation of 1000 consumers used a hill-climbing algorithm to discover the optimal price and subsequent profit for two common pricing strategies, pure bundling (where all articles are sold as a bundle) and linear (per-article) pricing for varying values of k_i . w_0 was held constant at 500 and k_i was varied between 0.001 and 1.0. A graph comparing k_i to optimal profit for each strategy is shown in figure 5.

There are several things to note regarding this experiment. The first is that the producer was able to learn the price which produced the optimal profit for both of these schedules using a hill-climbing algorithm. Since both of these schedules are one-parameter functions, we say that the producer successfully learned one characteristic of the congregation.

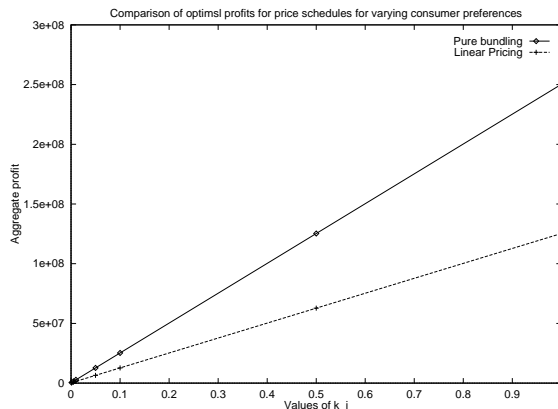


Figure 5: Results of Experiment 5 - a comparison of price schedules for varying consumer preferences.

Second, the experimental results shown here confirm the theoretical predictions in (Fay & MacKie-Mason 1998), namely that pure bundling dominates linear pricing in a steady state equilibrium. This is encouraging in that it indicates that further simulation-based experiments may be able to produce results that are in accordance with, and possibly contribute to, pricing theory.

Third, even in these simple experiments, these functions can contain discontinuities. As we introduce more complex environments containing multiple producers or consumers that behave strategically, we expect to see profit landscapes with multiple optima, necessitating more complex learning mechanisms. As in the previous experiment, simple techniques are sufficient for initial problems, but, as complications are added, we expect that these simple techniques will break down. However, as in the previous experiments, the simple mechanisms may be useful to bootstrap the problem, either as a means of reducing the search space for a more complex strategy or providing some initial guidance. In future work, we plan to deal with this problem using both neural networks trained using reinforcement learning and explicit modeling of consumers.

Conclusions

In this paper, we have argued that the formation of congregations is one possible way for agents to deal with the complexity of large-scale multiagent systems. By generalizing away from the specifics of how agents interact and what their goals are, we hope to create a general framework that can be applied to different sorts of multiagent domains.

We have suggested two classes of strategies for the formation of desirable congregations: the introduction of external mechanisms, such as taxes, and the introduction of learning to members of the agent population. So far, we have examined each strategy separately, so as to better understand its effects. A comprehensive solution will quite probably contain elements of both strategies.

A domain of particular interest to us is that of information economies. We believe that applying the principles and ideas of congregation formation to the creation of niches in the markets of information goods will allow designers of information-producing agents to better identify markets in which both they and their consumers can thrive.

Clearly, this work is still in its infancy. We have yet to integrate any of the dynamics that make real-life economies so difficult to understand. While simple solutions such as transaction fees and hill-climbing are sufficient for achieving excellent results in a simple model, as the system becomes more complicated, these methods produce cruder results. This, we believe, is where ontological information about both the consumers in an economy and the goods that are being produced can come into play. (Tesauro & Kephart 1998) indicate that learning can be used to reduce some of the dynamic fluctuations in an information economy, allowing producers and consumers to settle into a stable set of congregations. In their work, producers model the consumers as a profit signal which they then try to maximize. Consumer preferences are not modeled explicitly. Likewise, it is assumed that consumers do not reason strategically about producers. However, (Vidal & Durfee 1998) shows that there are situations in which it is beneficial for the producer to use a more detailed model of the consumer, either to model the consumer's reasoning about the producer or to exploit ontological information about the consumer's preferences.

This ontological information about consumer preferences and information goods may also be used in the same manner as an external transaction fee, providing a particular description of a congregation which can be used to attract or discourage potential consumers. This system-wide learning, in which the correctness of a label may depend on the labels chosen by other congregations, may be conducted in a distributed fashion or by an external agent, as was done in the SMS.

References

- Bakos, Y., and Brynjolfsson, E. 1998. Bundling information goods: Pricing, profits and efficiency. *Manufacturing Science*. also available at: <http://www.stern.nyu.edu/bakos/big.pdf>.
- Chuang, J. C.-I., and Sirbu, M. A. 1998. Network delivery of information goods: Optimal pricing of articles and subscriptions. Available at: <http://ksgwww.harvard.edu/iip/econ/chuang.html>.
- Durfee, E. H.; Mullen, T.; Park, S.; Weinstein, P.; and Vidal, J. 1998. *Cooperative Agents II*. Springer-Verlag. chapter The Dynamics of the UMDL Service Market Society.
- Fay, S. A., and MacKie-Mason, J. 1998. Competition between firms that bundle information goods. Unpublished.
- Hogg, T., and Huberman, B. A. 1991. Controlling

chaos in distributed systems. *IEEE Transactions on Systems, Man and Cybernetics* 21(6):1325–1332.

Kephart, J. O.; Hanson, J. E.; Levine, D. W.; Grosz, B. N.; Sairamesh, J.; Segal, R. B.; and White, S. R. 1998. Dynamics of an information-filtering economy. In *Proceedings of the Second International Workshop on Cooperative Information Agents (CIA)*.

Kephart, J. O.; Hanson, J. E.; and Sairamesh, J. 1998. Price and niche wars in a free-market economy of software agents. *Artificial Life* 4:1–23.

Tesauro, G. J., and Kephart, J. O. 1998. Foresight-based pricing algorithms in an economy of software agents. In *Proceedings of the First International Conference on Information and Computational Economics*.

Vidal, J. M., and Durfee, E. H. 1998. Learning nested agent models in an information economy. *Journal of Experimental and Theoretical Artificial Intelligence* 10.

Weinstein, P. C., and Birmingham, W. P. 1998. Agent communication with differentiated ontologies: Eight new measures of description compatibility. submitted for journal review.