

# Specialists, Generalists and Market Segmentation: A Multi-Agent Model <sup>†</sup>

Maksim Tsvetovat, Kathleen Carley, Katia Sycara  
maksim@cs.cmu.edu, kcarley@ece.cmu.edu, katia@cs.cmu.edu  
*Carnegie Mellon University*

September 10, 2001

**Abstract.** It has been observed in many instances that markets have a tendency to segment themselves into distinct sub-markets. This paper presents a multi-agent model that illustrates emergent market segmentation. The model illustrates the way local optimization processes result in an emergent global behavior.

## 1. Introduction

With emergence of electronic transactions as a predominant form of business-to-business trading (Fre00; Fas00; Che99), the study of such markets becomes more and more important. While in many ways electronic markets resemble the traditional trading floor-based marketplaces, the immediacy of available data on the behavior of such markets as a whole as well as individual market agents makes electronic markets an alluring subject of study.

It has been noted that unregulated markets have a tendency to become increasingly segmented over time (HF77).

A market is often described in terms of competition and natural selection. Organizations or individual may fail to flourish in certain environmental circumstances because others compete with them for essential resources. As long as the resources which sustain the market are finite and market participants have unlimited capacity to expand their business, competition must ensue.

A. Howley (How50) shows in his model that competition processes typically involve four components:

- demand for resources exceeds supply,
- selection eliminates weakest competitors
- competitors differentiate territorially or functionally, yielding a division of labor in a number of market niches,

---

<sup>†</sup> Tsvetovat, M., Carley, K. M., & Sycara, K., (2001), Specialists, Generalists and Market Segmentation: A Multi-Agent Model. Presented at *CASOS 2001 Conference*



- competitors within a niche become more similar as standard conditions of competition bring forth a uniform response.

Part of the efficiency resulting from increasing specialization is derived from the lower requirements for excess capacity. Given some uncertainty, most organizations maintain some excess capacity to insure the reliability of performance. In a rapidly changing environment, the definition of excess capacity is likely to change frequently. What is used today may become excess tomorrow, *visa versa*. Thus, a generalist organization (spanning more than one market niche), would be required to maintain excess capacity in every one of the covered niches, and becoming more specialized would reduce the excess capacity costs.

Other effects of agent specialization include lowered transaction costs and lower overall network load - both due to tightening of social networks within the market niches.

In this paper, we present a multi-agent model that demonstrates the emergence of agent specialization and market segmentation in an environment populated by self-interested agents. The model is used in a set of experiments that show that local optimization behavior of profit maximization leads to an emergent global optimization of the market via decrease in transaction costs and decrease in communication link loads.

## 2. Hypothesis

- A Market Populated With Self-Interested Agents Will Organize Itself Into Specialized Sub-Markets.
- Results of this specialization will be:
  - Decreased amount of communication needed to execute a transaction
  - Lower average transaction cost
  - Greater overall social welfare.

## 3. Market Representation

Let the market be represented as

$$M = \{G, p, A, C\}$$

where  $G$  is a set of goods  $g_i$  that are traded in the market,  $p$  is the market protocol,  $A$  is a set of agents and  $C$  is a set of connections between the agents.

The market protocol  $p$  is a set of finite state machines that describe the behaviors of agents participating in transactions based on their role. For example, an auction protocol specifies behaviors of the seller, the auctioneer and many bidders (who are identical except in their reservation prices). Such protocol would also include rules for winner determination and execution of the transaction.

The simulated market we used for our experiments consists of a central market clearing agent, and a number of trader agents (see Figure 3).

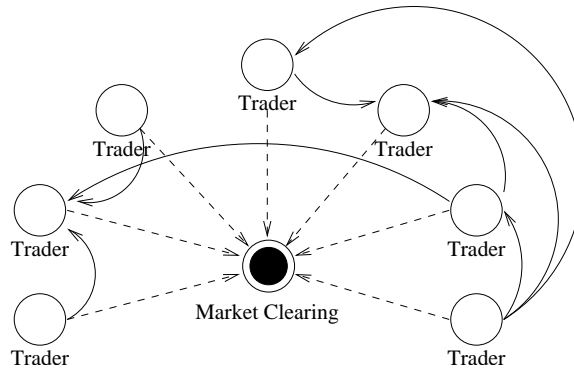


Figure 1. Simulated Market Structure

The trader agents within the market follow a Continuous Double Auction (CDA) protocol with periodic clearing (Bic99). In the CDA protocol, agents negotiate the transactions by submitting buy and sell bids to other agents. If an agreement is reached, the result of the transaction is reported to the auctioneer. The auctioneer collects transactions over a specified interval of time, then clears the market at the expiration of the bidding interval (Smi96).

The CDA protocol with periodic clearing was chosen for a number of reasons:

- Most commodity markets, as well as some of the stock markets, operate on protocol very close to CDA.
- The auctioneer, as a transaction-clearing entity, is an integral part of real-world commodity exchanges.
- While transaction clearing is centralized, negotiation is distributed.

- The periodic nature of the market clearing mechanism allows us to take snapshots of market activity and analyze the behavior of the market in discrete intervals.




Figure 2. A set of agents

Agent  $a_i \in A$  (see 3) represents a trader in the market and is defined as

$$a = \{BuyList, SellList, connections\}$$

where the *BuyList* and *SellList* are lists of goods that the agent deals with for buying and selling, with their reservation prices. The agent is free to add or drop goods from either of the lists as the agent's local strategy may dictate.

The agent maintains a list of *connections* which allows the agent to determine who it can deal with in regards to a particular good. A connection can be considered established between agents  $a_i$  and  $a_j$  with regard to good  $g_n$  if the agents have participated in a transaction governed by protocol  $p$  with regard to good  $g_n$ . Similarly, agents can establish new connections or drop connections at will.

The global set of connections  $C$  is the union of all sets of connections of all agents, and is the chief object of analysis in this paper.

### 3.1. UTILITY AND SELF-INTEREST

All agents in the system are designed to be self-interested and myopic. That is, in making a decision (such as “Should I go forward with this transaction” or “Should I add this item to my inventory”) the agents are only concerned with their own immediate profit. Agents have no way to estimate other agents’ profits or the global welfare of the market, or to predict the direction that the market will take in the future. The main goal of the agent is to execute the buy and sell orders it receives from its customers. Thus, such an agent is a fairly accurate representation of a market trader that specializes in negotiating and executing transactions but allows his or her customers to make their own buying decisions. The agent utility from each transaction is:

$$U = \textit{TransactionPrice} - \textit{ReservePrice} - \textit{TransactionCost}$$

where the transaction price is the final price at the end of the negotiations and the reserve price value has been supplied by the customer.

The transaction cost depends on the specific market and can be calculated in a number of different ways:

- *TransactionCost = const*: In this extreme case, the agents would seem to have no incentive to re-organize their communication networks, and will engage in as many communicative acts as it takes to maximize their utility. However, this case is unrealistic if we presume that the agents are subject to *bounded rationality*. In particular, processing one communicative act takes an agent a certain non-zero amount of time - and while the agents might be able to process a large number of messages, this number would still be finite during a finite clearing period. Thus, agents in the no-cost environment will still have an incentive to keep the number of messages low - with the lower bound set by a ratio of time it takes to process a message and length of the clearing period.
- *TransactionCost = f(m) where m is the number of messages*: In this case, the agents have a clear incentive to optimize their communication patterns with the goal of bringing the number of messages required to complete a transaction to a minimum. They can do that by creating a social network of other agents. Stronger ties in this network are created to agents that are most likely to possess the goods in question and sell them at a fair price.

### 3.2. AGENT DECISION-MAKING

When an order comes in, a trader agent executes the transaction if the good is on *buyist* or *sellist* and if the reservation prices supplied by the customer are acceptable in the current market conditions.

If an agent fails to complete the transaction within the clearing period, it does not receive the positive utility, but still has to pay the communication costs. Thus, it is possible that an agent completes some of the clearing periods with a negative utility.

The agents in the market have to make a set of decisions to complete the transaction. The fitness of these decisions to the market situation is largely responsible for whether an agent will be successful (receive high utility) or not.

In each clearing period, the agents must make the following decisions:

- Which of the market goods should be traded?
- Which agent should I talk to?
- Is the offer I received good enough?

The market goods are organized in a priority queue, sorted by the running-average utility derived from trading a particular good. At the beginning of each clearing period, the agent chooses an item from the list using the Metropolis criteria (i.e. an exponentially distributed random variable with  $\lambda$  proportional to the highest average utility among the goods).

The result of using the Metropolis criteria is that while the agent receives approximately equal average utilities from trading different goods, the chance of one of the goods is approximately equal. However, if one of the goods is more profitable than others, the chance of it being chosen increases dramatically. As the market prices vex and wane, the value of  $\lambda$  changes, and probabilities of other goods being picked increase.

The decision on which agent one should talk to is done in a similar manner. Each trader maintains a table where it stores a set of values  $U_g, a$ , a running-average utility of trading good  $g$  with agent  $a$ . The utilities for each of the goods are stored in a priority queue.

The negotiation proceeds as follows:

- First, an agent chooses a good to trade.
- Using the Metropolis criteria, it then chooses a potential partner and sends a request for bids.

- On receipt of the bid, the agent evaluates the price and chooses to accept or reject the bid.
- If the bid is rejected, the agent returns to step 2 and chooses another agent to talk to.
- If a bid is accepted, both agents report the transaction to the market clearing agent and record their respective utilities, thus updating the selection tables

#### 4. Accurate Transaction Modelling

In building a multi-agent model of a social phenomenon, one has to take into account several major factors.

First of all, the multi-agent model must accurately represent the processes present in the subject of study. In many cases it requires building a large knowledge base that mimics the cognitive processes of a human involved in the similar situation. An example of such cognitively accurate model is Soar (JWvLL95), which employs complex rule-based reasoning and learning processes to emulate performance and cognitive processes of a human operator.

However, in order to create accurate models of emergent phenomena, such as market activities, one must create simulation environments that include large numbers of agents. Due to their complexity and computational requirements, Soar and other cognitively accurate models cannot be used as part of large multi-agent system. Thus, emergent behaviors are often modelled using large collections of very simple agents (cellular automata). However, the simpler agents often cannot replicate the processes involved in market environments.

Thus, in order to build scalable and yet accurate simulation of a market, one must find a balance between cognitive complexity and computational requirements of a large number of agents.

We find that such a balance can be achieved with accurately modelled transaction protocols and realistic decision functions based on agent's self-interest.

#### 5. MarketSim - The Multi-Agent System

The MarketSim system was built on the foundation of the RETSINA multi-agent system framework. The RETSINA framework provides communication functionality (SDP<sup>+</sup>96), yellow and white-pages services (DSW96;

SLKW99a) that allow agents to find each other and interoperability (SLKW99b) services. The system allows wide re-use of tools and agents across different projects.

In building this model, we have used the communication and advertising components of the RETSINA architecture, a AgentFactory (ETSPar; ETSPar) protocol compiler and a logging and display services agent.

## 6. Analysis Tools

The RETSINA system was built with the presumption that behavior of multi-agent systems as a whole can be studied in conjunction with building of the agents on a micro level. Thus, the architecture provides for an easy way to integrate global logging and visualization into the multi-agent system.

The resulting logs are usually used to visualize the agent behavior via a tool called DemoDisplay. The DemoDisplay shows each agent as an icon on a field, and shows the communications between agents using animated arrows. When two agents communicate, their icons move closer to each other, and when they stop communicating the agents retreat to their initial positions around the border of the screen.

Thus, the screen becomes a real-time display of the social network of agents, showing clusters of agents as they communicate or conduct transactions.

The output of the logging tools can be also converted into MatLab data for further analysis. In this project, the logging data is converted into agent adjacency matrices and network diagrams. It is possible to sample the network diagrams at different times within the same system, so one can analyze the dynamics of relationships between agents.

## 7. Virtual Experiments with MarketSim

In the beginning of the paper we have stated a hypothesis that the despite being myopic and self-interested, the agents, in their quest for higher profits, will produce a market that is optimized on global scale. With the development of the MarketSim model, it is now high time to prove the validity of the hypothesis through a set of virtual experiments.

In these experiments, we have used several markets containing 25, 50 and 100 agents. All agents are similar in the way they compute their utilities, with the exception of reservation prices at every transaction. In all experiments, transactions are done via first-price English auction.

Agents are cognitively and conversationally limited - i.e. an agent can only participate in a small number of auctions at the same time due






Figure 3. Social Network for Baseline Market and its Adjacency Matrix

to both communication and reasoning load. There are 3 distinct goods to be traded. An agent starts out randomly initialized to be interested in 1 or more for both buying and selling.

The timekeeping for all agents is provided by a controlled clock, allowing the experimenter to slow down the agents to an observable speed, or increase the performance to near-realtime. In runs involving larger numbers of agents, it is often necessary to slow down the simulation clock to prevent overloading of the network interfaces.

The independent variables are:

- market size - the number of agents in the market
- market saturation - proportion of agents offering each of the goods; market saturation is equivalent to social network term "initial probabilistic density".

The main measurements in each of the experiments are:

- Average Network Load - the number of messages that pass through the network in one time period
- Average Transaction Cost - proportional to number of messages that have to be exchanged to complete one transaction
- Average Agent Utility - how well do the individual agents do?
- Overall Social Welfare - how well does the whole market do?

## 7.1. UNDIFFERENTIATED MARKETS - A BASELINE

Let us start with the most basic of the markets. Each of the market agents is randomly initialized with a set of goods to sell and a set of buy orders to fill.

In this run, the agents are not capable of changing their buy and sell lists, or estimating their future utility. They just blindly trade in auctions, and attempt to maximize their utility one auction at a time.

At the end of 20 auction periods, the messages were collected and collated into an adjacency matrix and a social network. The social network diagram on figure 6a illustrates the connections that traders had to make to establish to other traders in order to complete one transaction. Even given fairly advance matchmaking tools (SLKW99b), an agent had to communicate with as many as 12 other agents before being able to complete a transaction.

As the adjacency matrix (figure 6b) shows, there is no clear pattern to which agents have to communicate in order to complete transactions, which results in higher transaction costs and necessity for more communication to achieve a result.

Table I. Transaction Cost in Baseline Configuration

Market Size	Market Saturation		
	0.1	0.5	0.8
25	4.7200	18.0800	23.2000
50	9.4800	36.0400	47.4000
100	18.9800	74.5600	94.9200

Table II. Average Agent Utility (per Transaction) in Baseline Configuration

Market Size	Market Saturation		
	0.1	0.5	0.8
25	0.81	0.63	0.47
50	0.72	0.41	0.28
100	0.45	0.37	0.14

Table III. Average Network Load (msg/sec) in Baseline Configuration

Market Size	Market Saturation		
	0.1	0.5	0.8
25	7.2000	22.1000	28.9000
50	21.9000	92.8000	118.2000
100	92.2000	364.9000	473.4000

## 7.2. BASELINE RESULTS

The baseline results are not at all surprising, or even useful. There is a strong correlation between the average network load and market saturation - higher market saturation results in necessity for more communication to achieve a similar result. As market size and saturation increase, transaction costs go up, thus bringing down average agent utility.

The baseline experiment shows that a undifferentiated continuous double auction-based market is only viable while the number of agents involved is small and the market saturation is low. Any increase in the market saturation or market size results in rapidly increasing costs and, therefore, decreasing utility of individual agents as well as overall social welfare.

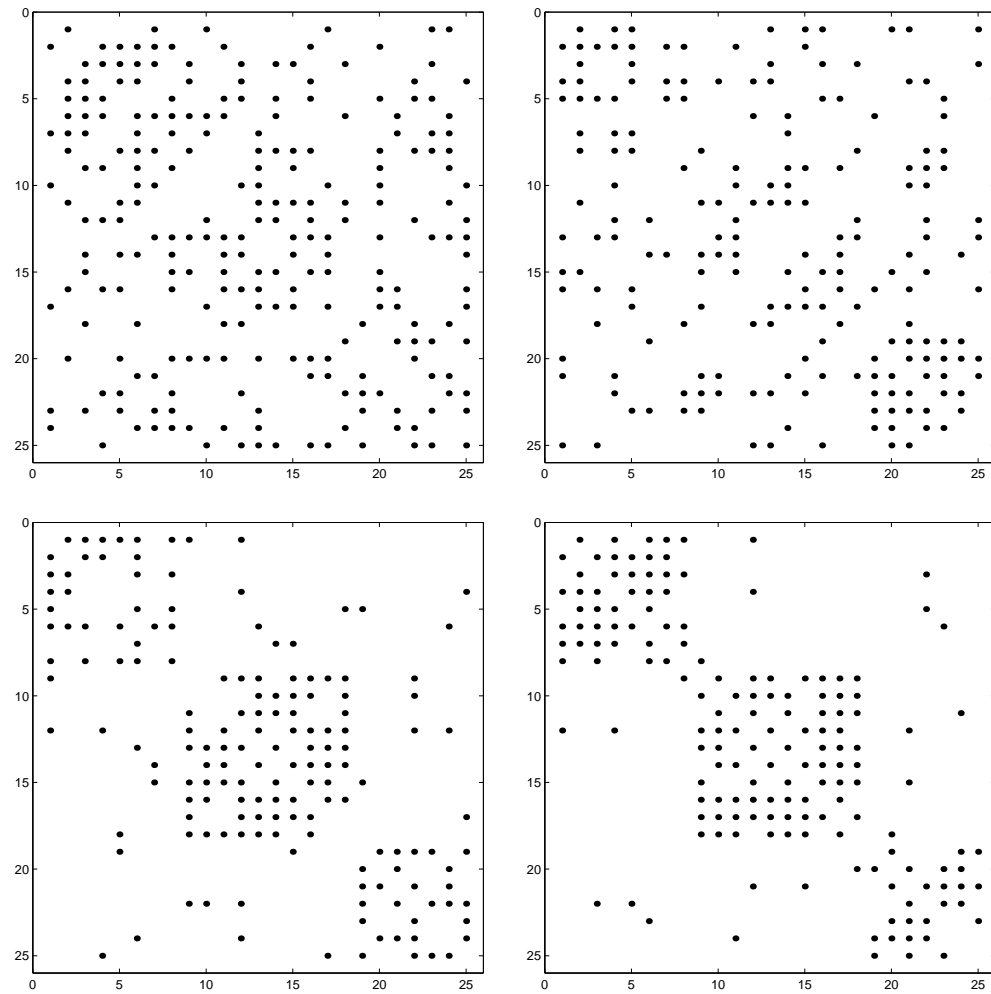
In a large and saturated market, there is a greater chance that lower-priced goods will appear, however the competition for them is very intense and many agents are forced to contend with buying higher-priced versions of the same good.

## 7.3. EMERGENCE OF MARKET SEGMENTATION

As Hannan (HF77) states, "...Organizations may insure reliable performance by creating specialized units..." or retreating into market niches that allow a highly specialized organization to thrive.

To simulate this process, agents were allowed to add and drop goods from their lists, based on the utility they gain from the transaction - thus allowing an agent to become as much of a generalist or specialist as the market conditions allow.

As the agent does business, utility from each transaction is normalized to be in 0-1 range, and running averages for each good are kept. The probability of the agent dropping an item from its list is directly proportional to the normalized utility.



*Figure 4.* Emergent market segmentation, 25 agents : (a) Unsegmented market, (b) After 100 clearing periods, (c) After 250 clearing periods, (d) After 500 clearing periods

Adding items is a risky proposition. In the real world, it might be possible to make estimates of what profits other agents are making, but in this simulation this data is purely private. Thus, adding goods becomes a matter of chance. The probability of an agent adding a good to its inventory is inversely proportional to the overall normalized utility (i.e. the lower the agent's profits, the more likely it is to try a new line of business).

The following set of adjacency matrices is the snap-shot of one of the runs of the system (Figures 7.2 and ??). Patterns similar to these were observed in all of the 20 runs of the market simulator.

#### 7.4. GLOBAL PATTERNS FROM LOCAL BEHAVIOR

The emergent specialization has a profound effect on the market conditions (see Figure 7.4). As agents specialize in selling one particular item, the network load decreases dramatically. As a consequence of a lower network load the transaction cost also decreases, which allows agents go get higher per-transaction utility.

Table IV. Effects of Market Specialization

	a	b	c	d
Network Load	21.20	18.28	12.44	8.80
Transaction Cost	18.2	13.41	9.3	4.2
Overall Saturation	0.3280	0.3312	0.2576	0.2320
Average Agent Utility	0.45	0.42	0.68	0.73

The overall market saturation also decreases, thus limiting the amount of competition in each of the market sectors and virtually eliminating any cross-talk between different sectors. This does not sound like good news for the market. However it has been noted in the literature (HF77) that a market shakedown often occurs after initial explosion. At the end of each shakedown the number of agents in a given market sector stabilizes at the maximum number of agents that can be sustained in the sector.

## 8. Emergent Social Networks

As seen from the adjacency matrix diagrams in figure 7.2, the agents' society quickly distributes itself into a set of densely interconnected clusters. Within the clusters, the networks are almost fully connected. For example, within cluster 1, consisting of 8 agents, the degree of interconnect (i.e. ratio of number of connections and number of possible connections) within the cluster is 6.5, and only 0.16 outside the cluster.

However, some agents evolve as gatekeepers(CITE gatekeeper). A gatekeeper is an agent .... blah blah blah

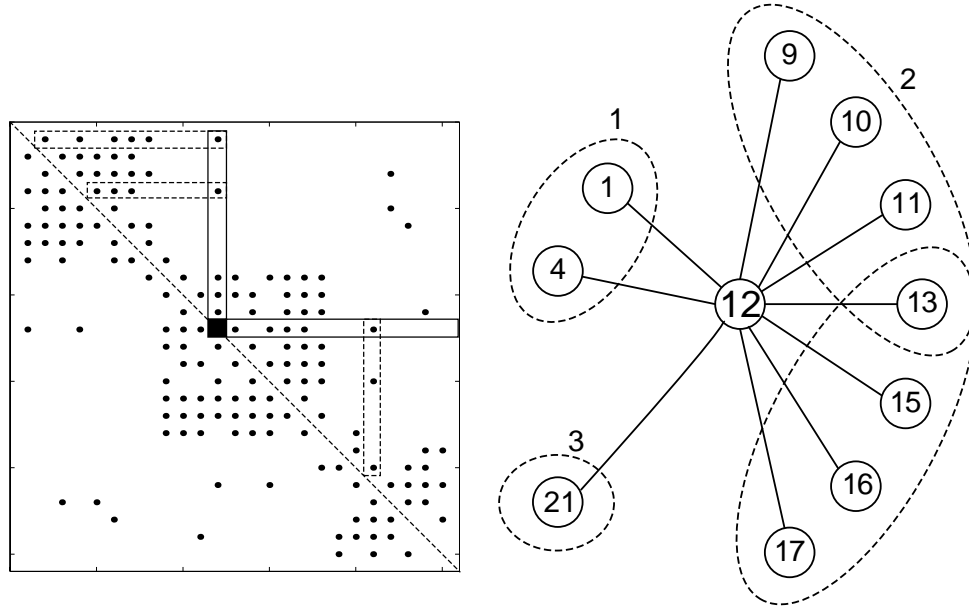


Figure 5. Agent 12, the Gatekeeper: (a) Position in the adjacency matrix, (b) Social Network graph

### 8.1. THE WORLD AND ILLUSTRIOUS LIFE OF AGENT 12

Let us examine what life would be for a gatekeeper agent within a segmented market. Agent 12 (see figure 8.1) is one of the most successful agents in the system,

## 9. Conclusions and Future Work

In this paper we have demonstrated a multi-agent model of a marketplace populated by self-interested adaptive agents. The model illustrates the segmentation of commodity markets by specialty - an emergent behavior borne out of local profit maximization motives. However, the local behaviors result in advancement of the global good - since the increase in segmentation of the market resulted in higher utility values, lower transaction costs and lower network loads for all agents in the market.

However, the model does not yet completely reflect patterns of interaction that occur in real markets. One of the most important aspects in terms of specialization is the advent of organizations of traders. In a real market, a trading firm employs many traders specializing in different

sectors of the market. However, the utility calculation and adaptation is done at the managerial level, which is above the market segments.

It is in our immediate plans to introduce management and hiring protocols and decision-making structures into the MarketSim - which would allow agents to hire each other and create organizations. The model will then be used to study emergence of organizational structures and effects of organizational structure upon the market performance of a firm.

Due to its distributed nature, the MarketSim model has proved to be a very scalable system, accommodating large numbers of agents and dense social networks. However, there remain several bottlenecks that limit the message throughput of the agents. In particular, in larger markets the blackboard was found to have delays. In the future, these bottlenecks will be eliminated so larger studies could be done to estimate the role of scale effects in simulated markets.

## References

- Martin Bichler. Decision analysis - a critical enabler for multi-attribute auctions. In *Global Networked Organisations, Proceedings of the 12 International Bled Electronic Commerce Conference*, June 1999.
- Chemdex.com, 1999.
- K. Decker, K. Sycara, and M. Williamson. Matchmaking and brokering. In *Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS-96)*, Dec 1996.
- G. Economou, M. Tsvetovat, K. Sycara, and M. Paolucci. Interaction without commitments: An initial approach. In *Proceedings of International Conference on Autonomous Agents*, 2001, to appear.
- FastParts.com, 2000.
- FreeMarkets.com, 2000.
- M. T. Hannan and J. Freeman. The population ecology of organizations. *American Journal of Sociology*, 82:5:929-964, 1977.
- A. H. Howley. *Human Ecology: A Theory of Community Structure*. Roland, 1950.
- R.M. Jones, R. Wray, M. van Lent, and J.E. Laird. Planning in the tactical air domain. In *Planning and learning: On to real applications, papers from the 1994 Fall symposium (Technical Report No. FS-94-01)*, Menlo Park, CA, 1995. AAAI Press. Soar 95.27.
- K. Sycara, K. Decker, A. Pannu, M. Williamson, and D. Zeng. Distributed intelligent agents. *IEEE Expert - Intelligent Systems and Their Applications*, 11(6):36-45, 1996.
- K. Sycara, J. Lu, M. Klusch, and S. Widoff. Dynamic service matchmaking among agents in open information environments. *ACM SIGMOD Record, Special Issue on Semantic Interoperability in Global Information Systems*, 1999.
- K. Sycara, J. Lu, M. Klusch, and S. Widoff. Matchmaking among heterogeneous agents in the internet. In *Proceedings AAAI Spring Symposium on Intelligent Agents in Cyberspace*, 1999.

Vernon Smith. Economics in the laboratory. Technical report, University of Arizona, 1996.