# **Electronic Market Making: Initial Investigation**

#### Yuriy Nevmyvaka

The Robotics Institute Carnegie Mellon University Pittsburgh, PA 15213 yuriy@cs.cmu.edu

### Katia Sycara

The Robotics Institute Carnegie Mellon University Pittsburgh, PA 15213 katia@cs.cmu.edu

### Duane J. Seppi

GSIA Carnegie Mellon University Pittsburgh, PA 15213 ds64@andrew.cmu.edu

#### **Abstract**

This paper establishes an analytical foundation for electronic market making. Creating an automated securities dealer is a challenging task with important theoretical and practical implications. Our main interest is a normative automation of the market maker's activities, as opposed to explanatory modeling of human traders, which was the primary concern of earlier work in this domain. We use a simple class of "non-predictive" trading strategies to highlight the fundamental issues. These strategies have a theoretical foundation behind them and serve as a showcase for the decisions to be addressed: depth of quote, quote positioning, timing of updates, inventory management, and others. We examine the impact of various parameters on the market maker's performance. Although we conclude that such elementary strategies do not solve the problem completely, we are able to identify the areas that need to be addressed with more advanced tools. We hope that this paper can serve as a first step in rigorous examination of the dealer's activities, and will be useful in disciplines outside of Finance, such as Agents, Robotics, and E-Commerce.

#### 1 Introduction

What is market making? In modern financial markets, market makers (or dealers) are agents who stand ready to buy and sell securities. The rest of market participants are therefore guaranteed to always have a counterparty for their transactions. This renders markets more orderly and prices less volatile. Market maker are remunerated for their services by being able to "buy low and sell high". Instead of a single price at which any trade can occur, dealers quote two prices – a "bid" (dealer's purchase, customer's sale) and an "ask" (dealer's sale, customer's purchase). The ask is higher than the bid, and the difference between the two is called the spread – the dealer's source of revenue.

What are the benefits of automating this activity? This is a challenging decision problem: how can a machine update the bid-ask spread, anticipating or reacting to changes in the supply and demand for a security? This

setting is also a great test bed for Machine Learning and Statistical techniques. Creation of an electronic dealer is a stab at the main goal of AI: replication of a human decision process, which is notoriously difficult to model or imitate. From a more pragmatic point of view, electronic market makers could render securities professionals more productive and markets more stable. Automated dealers, if designed properly, will not engage in market manipulations and other securities laws violations that recently resulted in a number of dealer-centered scandals in both the NASDAO [3] and NYSE [9]. Also, a more in-depth understanding of the dealer's behavior will give us better guidance in extreme situations (market crises) and will facilitate the regulatory oversight. Finally, we expect automated market making to have an impact on other disciplines that employ various market mechanisms to solve distributed problems: Robotics [15], E-Commerce [7], Intelligent Agents [12], etc.

The goal of this paper is to establish an analytical framework for electronic market making, using a simple class of strategies to highlight some central issues and challenges in this domain. The paper is organized as follows. Section 2 explains where the present effort is situated relatively to other research in this area. In Section 3, we describe our experimental setup. Section 4 presents a simple model of electronic market making and a general taxonomy of possible strategies. Section 5 makes a case for so-called "non-predictive" market-making strategies, while Section 6 presents the relevant experimental results. We conclude with a recap of important issues and a description of future work.

## 2 Related Work Comparison

Automation of dealer's activities was suggested more than three decades ago [2] and is an important part of Market Microstructure – an area that has evolved into an independent subfield of Finance. The bulk of previous research on market making is mostly concerned with the sources and components of the bidask spread. A number of models have been developed to explain the evolution of the spread, incorporating various factors that affect the market maker's decision process, such as inventory [13], information [4], volatility [6], risk aversion, competition [14], and many others [8]. The problem with such approaches is that

they are mostly explanatory in nature. The relevant work in the Computer Science community is more limited. Market making has been adopted as a test-bed for new Machine Learning techniques [11] with a goal to demonstrate the general effectiveness of a learning algorithm, as opposed to treating market making as a problem that requires solving. Also, empirical work has demonstrated the limitations of hard-coding marketmaking "rules" into an algorithm [5]. The motivation behind our approach is fundamentally different from that of previous research. Since we are interested in creating an electronic market maker, we are much more concerned about future performance. Therefore, our primary goal is to optimally change the spread over the next iteration instead of finding the best model for past transactions. We are trying to create something much more normative (as opposed to explanatory): to determine which factors are important for making the spread update, and to capture the decision process of a dealer.

### 3 Experimental Setup

In our experiments, we used the Penn Exchange Simulator (PXS) – software developed at the University of Pennsylvania, which merges actual orders from the Island electronic market with artificial orders generated by electronic trading agents [10]. Island is what is called an Electronic Communication Network (ECN). ECNs are somewhat different from traditional stock exchanges such as NYSE or the NASDAQ OTC market. NYSE and NASDAQ employ securities dealers to provide liquidity and maintain orderly markets, and use both market and limit orders. A market order is an instruction from a client to the dealer to buy or sell a certain quantity of stock at the best available price, whereas a limit order asks for a transaction at a specified or more advantageous price. Therefore, market orders guarantee the execution, but not the price at which transaction will occur, whereas limit orders guarantee a certain price, but transaction may never happen. Island ECN is a purely electronic market, which only uses limit orders and employs no designated middlemen. All liquidity comes from customers' limit orders that are arranged in order books (essentially two priority queues ordered by price) as shown in Figure 1a (limit price – number of shares).

If a new order arrives, and there are no orders on the opposite side of the market that can satisfy the limit price, then the order is being entered into the book. In Figure 1b, a new buy order for 1000 shares at \$25.20 or less has arrived, but the best sell order is for \$25.30 or more; thus no transaction is possible, and the new order is entered into the buy queue. When another buy order arrives for 250 shares at \$25.40 or less, it gets transacted (or crossed) with the outstanding orders in the sell queue: 150 shares are bought at \$25.30 and another 100 shares are bought at \$25.35 (Figure 1c). This demonstrates that even though there are no designated market orders in ECNs, immediate and

guaranteed execution is still possible by specifying a limit price that falls inside the opposite order book. All crossing is performed by a computer respecting the price and time priority, without any intermediaries.

25.56 - 300 25.55 - 1000 25.35 - 200 25.30 - 150	Sell Orders	25.56 - 300 25.55 - 1000 25.35 - 200 25.30 - 150	25.56 - 300 25.55 - 1000 25.35 - 100
25.21 - 200 25.19 - 300 25.15 - 785 25.10 - 170	Buy Orders	25.21 – 200 25.20 – 1000 25.19 – 300 25.15 – 785 25.10 – 170	25.21 – 200 25.20 – 1000 25.19 – 300 25.15 – 785 25.10 – 170
	(a)	(b)	(c)

Figure 1.

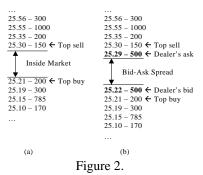
What PXS does is simple: at each iteration, it retrieves a snapshot of the Island's order book, gathers all of the limit orders from trading agents in the simulation, and then merges all the orders (real and artificial) according to the ECN rules described above: some orders transact and some get entered into the book. When transactions happen, agents are notified about the changes in their inventory and cash, and the new merged order book becomes available to all the agents to study and make decisions. This new order book is the state representation of the simulator's market, which can be different from the Island market because the orders from electronic traders are present only in the simulator. The inherent problem with such setup is that Island (real-world) traders will not react to the actions of the traders in the simulator, which can lead to a disconnect between the two markets. This implies that in order for the experiment to remain meaningful, the simulator traders have to remain "low impact" - i.e. their actions should not move the simulated price significantly away from the Island price. We enforce this property by prohibiting the participating agents from accumulating a position in excess of 100,000 shares either short or long. Such a simple rule gets the job done surprisingly well. To put thing in perspective, daily volume in the simulator reaches many million shares (actively traded MSFT is being used).

As stated before, PXS does not have market orders that flow through the dealers, or any designated dealers at all, for that matter. This can lead to a conclusion that such setup is ill-suited for studying the market maker's behavior. But we have to draw a distinction between market making as an institution (as seen on the NYSE floor) vs. market making as a trading strategy (used on proprietary trading desks and certain OTC dealing operations). The former can be interpreted as public service, where the market maker has certain obligations. He is supposed to be compensated by the bid-ask spread, but because of heavy regulations that protect customers, the dealer often finds himself being restricted in trading opportunities, which limits his profits [9]. Alternatively, market making can be

interpreted as a strategy where a trader tries to keep his stock position around zero (being "market neutral") and to profit from short-term price fluctuations. As far as low profile trading goes, the market maker is not supposed to "move" markets: the NYSE dealers are explicitly prohibited from doing so by the "negative obligation" principle. Thus, our setup is well suited for studying market making as a strategy, which also happens to be the main part of market making as an institution.

### 4 Market Making: A Model

We decompose the problem facing the electronic market maker into two components: establishing the bid-ask spread and updating it. We further subdivide the update methods into predictive and non-predictive. The primary objective of a dealer is to manage the bid-ask spread: it has to be positioned in such a way that trades occur at the bid as often as at the ask, thus allowing the dealer to "buy low and sell high". (We will examine these mechanics in Section 5). In order for this to happen, the quotes have to straddle the "true price" of the security [6] and be positioned as close to it as possible. However, the "true price" is an elusive concept, difficult to determine or model. Therefore, the first decision for the market maker (either human or artificial) is where to establish the initial spread.



There are two ways to approach this decision. The first, hard way is to perform the actual valuation of the security being traded: for a stock, try to determine the value of the company using discounted cash flows, ratios, etc.; for a bond, find the present value of the promised payments, and so on. If there is no established market, or the market is very illiquid, then valuation may be the only approach. Fortunately, the majority of modern securities markets employ limit orders in some capacity. The two queues of the order book should be an accurate representation of the current supply (sell queue) and demand (buy queue) for the security. Presented with such supply-demand schedule, the market maker tries to determine the consensual value. In the simplest case, the dealer can observe the top of each book – the best (highest) buy and the best (lowest) sell - also known as the "inside market". He then assumes that the market's consensus about the price lies somewhere between these two numbers. In Figure 2a the best bid is \$25.21, and the best ask is \$25.30 (the inside market is \$25.21-30), and the "true price" of the stock is in this interval. Now, the market maker can use the top of each book as a reference point for positioning his initial quotes – at \$25.22-29, for example (Figure 2b) – and then update his spread as the book evolves with new arrivals, transactions and cancellations.

Updating the spread is at the heart of market making. While the order book is informative about the consensus price of the security, it often fails to provide sufficient liquidity, thus creating demand for market makers. We classify update strategies into two categories. The first attempts to foresee the upcoming market movements (either from the order book misbalances or from short-term patterns), and adjust the spread according to these expectations. The second group reasons solely on the information about the current inside market. These "non-predictive" strategies are inherently simpler, and, therefore, better suited for our introductory examination of electronic market making.

### 5 Non-Predictive Strategies

In order to make the case that the non-predictive strategies are worth considering, let's examine in detail how the market maker earns money. Following the movement of the stock price over several hours, it is easy to discern some patterns: going up, down, back up, and so on. But if we take an extremely short time period (seconds or fractions of a second), it becomes apparent that the stock constantly oscillates up and down around a more persistent (longer-term) movement. If the price rises consistently over an hour, it doesn't mean that everyone is buying; selling is going on as well, and the transaction price (along with the inside market) moves down as well as up. Figure 3 illustrates this: while there is an upward movement (the dotted line), we can see the temporal evolution of the order book where transactions happen at the top of the buy queue, then the sell queue, then buy, then sell again. By maintaining his quotes on both sides of the market, at or close to the top of each order book, the market maker can expect to get "hit" at his bid roughly as often as at the ask because of these fluctuations. This way, after buying at the bid (low) and selling at the ask (high), the dealer receives the profit equal to the bid-ask spread for the two trades, or half-the-spread per trade. In the context of Figure 3, suppose that the top order in each queue is the dealer's; the dealer buys at \$25.10, then sells at \$25.18 (8 cents per share profit), buys at \$25.16 and sells at \$25.26 (10 cents profit). If each transaction involves 1,000 shares, and all this happens over several seconds, then market making can be quite profitable.

Having understood the nature of the dealer's income, we can re-formulate his task: adjust the bid-ask spread in such a way that the orders generated by other market participants will transact with the dealer's bid quote and the dealer's ask quote with the same frequency. In our

example, we are looking for an algorithm to maintain the dealer's quotes on top of each queue to capture all incoming transactions. The stock price is steadily going up overall, while fluctuating around this general climb, so if the market maker wants to maintain profitability, then his spread should also continuously move up straddling the stock price.



How can the dealer tell at any given point looking forward that it's time to move the spread up and by how much? The non-predictive family of electronic trading strategies would argue that he cannot and need not do so. It postulates that while there are some patterns globally, the local evolution of the stock price is a random walk. If this random walk hits the bid roughly as often as it hits the ask, then the market maker makes a profit. This means that an uptick in the stock price is as likely to be followed by a downtick as by another uptick - idea of efficient markets with short-term liquidity imbalances. If the above assumption holds, and if the dealer is able to operate quickly enough, then the trading strategy is very simple. All the dealer has to do is maintain his bid and ask quotes symmetrically distant from the top of each book. As the book evolves, the market maker has to revise his quotes as quickly as possible, reacting to changes in such a way that profitability is maintained.

In principle, the dealer should be market neutral – i.e. he doesn't care what direction the market is headed – he is only concerned about booking the spread. On the other hand, the dealer is interested in knowing how the inside market will change over the next iteration in order to update his quotes correctly. The way the non-predictive strategies address this is by assuming that the inside market after one short time step will remain roughly at the same level (that's the best guess we can make). Therefore, being "one step behind" the market is good enough if one can react quickly to the changes. Such is the theory behind this class of strategies, but in practice this turns out to be more complicated.

# **6** Implementation and Results

Here is a general outline of an algorithm that implements the non-predictive strategy; at each iteration:

- (1) Retrieve the updated order book;
- (2) Locate an inside market;

- (3) Submit new quotes (buy and sell limit orders), positioned relatively to the inside market;
- (4) Cancel previous quote.

As it should be clear from this description and the theoretical discussion above, there are three main factors, or parameters, that determine a non-predictive strategy: position of the quote relative to the inside market, depth of the quote (number of shares in the limit order), and the time between quote updates.

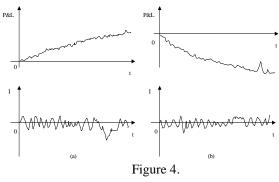
#### 6.1. Timing

Timing is, perhaps, the simplest out of the three parameters to address. Our experimental results are consistent with the theoretical model from Section 5: faster updates translate into higher profits. The dealer wants to respond to changes in the market as soon as possible, and therefore, the time between the updates should be as close to zero as the system allows. The implication of this rule is that update time should certainly be minimized when designing an electronic market maker. The computational cycle must be performed as fast as possible, and the communication between the dealer and the market should also be sped up. Our experiment shows that it's the latter issue that is a more important bottleneck, plus it's the one harder to control. The computational cycle for a simple strategy is inherently short (under 1 second), but it takes about 3 - 5 seconds for a submitted order to show up in the book, and about the same delay for the order to get cancelled (if not transacted) after it appears in the book. While these delays are not unreasonable by the real world's standards, they are not negligible. This is one of the "frictions" in market implementation, which should not be overlooked: the dealer wants to access the market as quickly as possible, but such delays can prevent him from operating on a scale short enough to capture the small fluctuations. Therefore, other systems where these delays can be decreased can potentially be more effective and produce better results than our simulated setup.

#### **6.2. Quote Positioning**

Position of the quote relative to the rest of the order book is the most important parameter. We use a simple distance metric: number of cents by which the dealer's quote differs from the top [non-dealer] order in the appropriate book. We started our implementation with a well-known, albeit controversial practice of "penny jumping". In general, penny jumping occurs when a dealer, after entering his customer's order into the order book, submits his own order, which improves the customer's limit price by a very small amount. The dealer effectively "steps in front" of his customer: the customer's potential counterparty will transact with the dealer instead. Such practice is not illegal (because the dealer does provide a price improvement over the original order), but is considered unethical, and became the center of the recent NYSE investigation [9]. In our case, we are simply undercutting the current inside market (or the "de facto"

bid-ask spread) by one cent on both sides. Figure 2a shows that if the inside market is 25.21-30, our market maker's orders will make it 20.22-29 (the size of the bid-ask spread goes from 9 to 7 cents). This way, the dealer is guaranteed to participate in any incoming transaction up to the size specified in the depth of his quote. We expect the following behavior from this strategy: the revenue (P&L) should rise slowly over time (since profit per share is tiny), while the inventory (I) ought to fluctuate around zero (see Figure 4a). We observe, however, that in our test set a typical trading day looks like Figure 4b: the strategy gradually loses money.



The fundamental problem is that while we base our decision on the book at time t<sub>0</sub>, our orders gets placed in the book at t<sub>1</sub>, which may or may not be different from the original t<sub>0</sub> book. The non-predictive strategy places an implicit bet that the two books will be close - at least enough to preserve the profitable property of dealing. What actually happens in our tests, is that the inside market is already tight, plus the book changes somewhat over the 3 second delay, and thus, oftentimes, both the bid and the ask placed at to end up on the same side of the market at t<sub>1</sub> (Figure 5a). Then one of the orders transacts, and the other gets buried deep in the order book. If we find ourselves in this situation on a regular basis, we end up paying the spread instead of profiting from it. This explains why our actual P&L pattern mirrors the pattern we expected. This experiment highlights the three challenges for electronic market making in any setting: (1) making decisions in one book, acting in another; (2) market "frictions" aggravate this disconnect; and, (3) spreads are extremely tight leaving little or no profit margin [1].

Does this mean that the non-predictive strategies are inherently money-losing? Not at all – one modification solves the third issue and mitigates the first two. We found that instead of undercutting the inside market, it is more profitable to put the quotes deeper in their respective books. The dealer's spread is wider now resulting in higher margins, plus even if the quotes get put into the book with a delay, they still manage to straddle the inside market and preserve the "buy low, sell high" property. Figure 5b shows the exact same scenario as Figure 5a, but with wider dealer quotes. Wider spreads lead to higher profit margins, but less volume flowing through the dealer. One has to find a

balance between these two components of the revenue. We determined that putting the quotes 1-3 cents away from the inside market works well, and alleviates the concerns that make penny jumping unprofitable. We also determined that the third fundamental parameter – the depth of quote – can also increase the dealer's volume and thus his profitability: the deeper the quote the more stocks will flow through the dealer with each transaction. However, increasing the volume to an arbitrary level has consequences for dealer's inventory.

25.56 - 300	25.56 - 300	25.56 - 300	25.56 - 300	
25.55 - 1000	25.55 - 1000	25.55 - 1000	25.55 - 1000	
25.35 - 200	25.35 - 200	25.35 - 200	25.35 - 200	
25.30 - 150	25.33 - 500	25.33 - 500	25.33 - 500	
25.29 - 500	25.30 - 150	25.30 - 150	25.30 - 150	
	25.25 - 100		25.25 - 100	
25.23 - 500	25.24 - 300		25.24 - 300	
25.26 – 200				
25.22-1000		25.26 - 200		
25.19 - 300	25.22 - 750	25.23 - 500	25.23 - 250	
25.15 - 785	25.19 - 300	25.22-1000	25.22 - 1000	
25.10 - 170	25.15 - 785	25.19 - 300	25.19 - 300	
	25.10 - 170	25.15 - 785	25.15 - 785	
		25.10 - 170	25.10 - 170	
t0	t1	tO	t1	
(a)	ı	(1	o)	
Figure 5.				

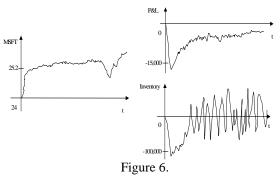
#### **6.3. Inventory Management**

In theory, the market maker should buy roughly as frequently as he sells, which implies that his inventory should fluctuate around zero. In practice, however, this doesn't always work out. If a stock price is going up consistently for some period of time, the dealer's ask ends up getting hit more often than his bid. If the dealer's quote is deep and/or close to the inside market he winds up with a growing short position in a rising stock - he is taking a loss. Plus, even the main assumption behind the non-predictive strategies can fail: when a stock "crashes", there are actually no buyers in the marketplace, and our entire marketmaking model is simply not valid any more. Finally, if a dealer accumulates a large position in a stock, he becomes vulnerable to abrupt shifts in supply and demand – i.e. if he has a significant long position, and the stock price suddenly falls, then he's taking a loss. In brief, there is a trade-off: on one hand, the dealer wants to have a large inventory to move back and forth from one side of the market to the other making profit, but then he doesn't want to become exposed by having a position that cannot be easily liquidated or reversed. To reconcile these conflicting goals, some rules have to be put in place to manage the dealer's holdings. We have implemented and tested a number of such approaches. The distance to the inside market proved to be the more effective, since it's tied both to the inventory and profitability. If there is too much buying (the dealer's ask is being hit too often, and he accumulates a short position), then moving the ask deeper into the sell book compensates for this. Also, if the stock is going in one direction consistently, this approach will force the spread to be continuously adjusted, using the inventory

misbalance as a signal. We achieved good results with the formula QuoteDistance = MinimumDistance + alpha \* max(0, Inventory - InitialLimit)/Inventory \* MinimumDistance. When the position is within the InitialLimit, the quote is always set MinimumDistance away from the market, but if the inventory gets outside the limit, the quote moves further away, encouraging the inventory's movement in the opposite direction.

### 7 Analysis and Conclusion

Implementing and testing the non-predictive marketmaking strategies, we arrived at a number of conclusions: faster updates allow to follow the market more closely and increase profitability; to combat narrow spread and time delays, we can put the quote deeper into the book, although at the expense of the trading volume; trading volume can be increased with deeper quotes; inventory can be managed effectively by resizing the spread. However, non-predictive strategies do not solve the market-making problem completely. Figure 6 exemplifies one general shortcoming of nonpredictive strategies: at the open, the price keeps going up, the market maker cannot get his quotes out of the way fast enough, accumulates a large short position, and loses a lot of money. All this happens in 10 minutes. This goes back to one fundamental problem: there are times when the short-term fluctuations in which the non-predictive strategies are rooted just aren't there. The only way to address this is to use some predictive instruments - order book misbalances, past patterns, or both - in order to be prepared for these streaks.



Even with this inherent weakness, the non-predictive strategies have some clear practical advantages. First, they are simple and computationally cheap, but, at the same time, a human trader can never replicate them. Their performance can be improved by speeding up the access to the market, or by applying them to less liquid stocks. Their use of the inside market as the only decision anchor makes them indifferent about the composition of the "trading crowd". And, finally, the problematic situations, described above can be handled by special cases to boost the overall performance.

In this paper, we have presented a structured framework for reasoning about the electronic market making and analyzed a number of fundamental issues in this domain using a simple class of strategies as an example. While we have not provided all the answers, our main goal was to frame electronic market making as a coherent problem and to highlight the points that must be addressed in order for this problem to be solved. We believe that this is an interesting and promising area, and that advances in the electronic market making will be useful in disciplines beyond Finance.

### **Bibliography**

- 1. Barclay M., Christie W., Harris, J., Kandel, E., and Schultz, P., *The Effects of Market Reform on the Trading Costs and Depth of NASDAQ Stocks*, Journal of Finance 54, 1999.
- 2. Black F., *Toward a Fully Automated Exchange*, Financial Analysts Journal, Nov.-Dec. 1971.
- 3. Christie, W. and Schultz, P., Why do NASDAQ market makers avoid odd-eighth quotes? Journal of Finance 49, 1813-1840, 1994.
- 4. Glosten, L., and Milgrom, P., Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders, Journal of Financial Economics 14, 1985.
- 5. Hakansson, N.H., Beja, A., and Kale, J., On the Feasibility of Automated Market Making by a Programmed Specialist, The Journal of Finance, Vol. 40, March 1985.
- Ho, T. and Stoll, H., Optimal Dealer Pricing Under Transactions and Return Uncertainty, Journal of Financial Economics 9, 1981.
- 7. Huang, P., Scheller-Wolf, A., and Sycara, K., *Design of a Multi-unit Double Auction E-market*, Computational Intelligence, vol. 18, No. 4, 2002.
- 8. Huang, R., Stoll, H., *The Components of the Bid-Ask Spread: A General Approach*, The Review of Financial Studies, Vol. 10, Winter 1997.
- 9. Ip, G. and Craig, S., NYSE's 'Specialist' Probe Puts Precious Asset at Risk: Trust, The Wall Street Journal, April 18, 2003.
- 10. Kearns, M. and Ortiz, L., *The Penn-Lehman Automated Trading Project*, IJCAI TADA Workshop, 2003.
- 11. Kim, A.J., Shelton, C.R., *Modeling Stock Order Flows and Learning Market-Making from Data*, Technical Report CBCL Paper #217/AI Memo #2002-009, M.I.T., Cambridge, MA, June 2002.
- 12. Klusch, M. and Sycara, K., *Brokering and Matchmaking for Coordination of Agent Societies: A Survey*. In Coordination of Internet Agents, A. Omicini et al. (eds.), Springer, 2001.
- 13. Madhavan, A., and Smidt, S., An Analysis of Changes in Specialist Inventories and Quotations, The Journal of Finance, Vol. 48, Dec 1993.
- 14. Seppi, D., Liquidity Provision with Limit Orders and Strategic Specialist, The Review of Financial Studies, Vol. 10, Issue 1, 1997.
- 15. Zlot, R.M., Stentz, A., Dias, M.B., and Thayer, S., *Multi-Robot Exploration Controlled by a Market Economy*, IEEE *International* Conference on Robotics and Automation, May 2002.