

# Learning, Evolution and Tick Size Effects in a Simulation of the Nasdaq Stock Market<sup>\*</sup>

Vince Darley<sup>a</sup>, Alexander Outkin<sup>b</sup>, Tony Plate<sup>c</sup>, Frank Gao<sup>d</sup>

BiosGroup, Inc.<sup>\*\*</sup>

March 31, 2001

## Abstract

This paper presents the results of our research into the behavior of a dealer-mediated stock market, similar to Nasdaq, by using an agent-based model of the market. We modeled on an individual level the decision-making process of market makers (dealers) and investors, as well as explicitly modeling market infrastructure and rules. The model allows investigation of market behaviors under a variety of scenarios and conditions. As our main goal, we investigated possible effects of tick size reduction, and found that in the simulated market environment it may result in decreasing the market's ability to perform the function of price discovery. Calibrating the model, we discovered that it exhibits a number of behaviors normally associated with real-world scenarios, such as the presence of fat tails, spread clustering, etc. We also created learning market makers and investigated their behavior and the strategies they use. We found that there is a variety of conditions under which artificial learning strategies outperform those extracted from the data or from the expert knowledge.

## Introduction

We have built an agent-based simulation of the stock market that represents a highly realistic picture of a dealer-mediated market, with the flexibility to model many features of real-world markets. In our research we concentrated on the following topics:

1. Investigating, mainly in a qualitative fashion, the consequences of regulatory and structural changes to the market (the most important being the question of minimum tick size).

2. Investigating whether our model is able to replicate some of the observed features of real-world markets.
3. Validating the model (this encompasses and includes the previous two points).
4. Designing learning agents, and investigating the behaviors they learn and their ability to perform profitably in the market.

Our results are significant in two respects. First, the model is robust in that the simulated market, investors, and market makers perform realistically under a wide variety of conditions. Second, the market dynamics produced by the model have the same qualitative properties as those observed in real markets. Thus the model provides a test bed in which to investigate the effects of changes in market rules and conditions. For example, we have derived results pertaining to volatility, liquidity, spread sizes, and spread clustering, with our main focus being on the market impact of reducing the tick size.

For the specific issue of minimum tick size, our results show that the market's ability to perform price discovery may be significantly impeded by reductions in tick size (for example, from \$1/16 to \$1/100) when the market contains individuals who utilize "parasitic" strategies (which, we can argue, correspond to the behavior of some day traders).

One important feature of our results is that often they are not a consequence of a uniquely identifiable feature of the model, or of the actions of certain market participants. Rather, they result from a relatively complex set of interactions of market makers, investors, market rules, and market infrastructure. Thus, even in a relatively simple setting, we can observe unintended consequences of the market's design.

---

<sup>\*</sup> The authors are grateful to Stuart Kauffman, Bennett Levitan, Richard Palmer, Isaac Saias, and others for their valuable comments and suggestions. We thank Paula Lozar for editing help. We acknowledge Nasdaq's financial support in working on this project. Usual disclaimer applies.

<sup>a</sup> Tel: 44 (0) 20 7420 4320; e-mail: [vince.darley@eurobios.com](mailto:vince.darley@eurobios.com)

<sup>b</sup> Corresponding author. Tel: (505) 992-6743; e-mail: [alexander.outkin@biosgroup.com](mailto:alexander.outkin@biosgroup.com).

<sup>c</sup> Tel: (505) 992-6712; e-mail: [tony.plate@biosgroup.com](mailto:tony.plate@biosgroup.com)

<sup>d</sup> Tel: (505) 992-6739; e-mail: [frank.gao@biosgroup.com](mailto:frank.gao@biosgroup.com)

<sup>\*\*</sup> BiosGroup, Inc. Phone 505-992-6700  
317 Paseo de Peralta Fax 505-988-2229  
Santa Fe, NM 87501 <http://www.biosgroup.com>

## Model Overview

We built an agent-based model of the stock market. The main actors in the model are market makers and investors. We model them as separate, autonomous entities whose interactions within the marketplace produce price discovery and determine the market's dynamics.

The market contains a single security whose fundamental value (also referred to as "true value") is exogenously specified: the underlying value of the security is assumed to fluctuate according to a random process.

Investors receive a "noisy" version of that fundamental value and act on the basis of that information; thus the fundamental value in the simulation can be interpreted as an aggregate source of quasi-information. There are investor types in the simulation who make their decisions based solely on the perceived fundamental value of the security, as well as those who follow trends, etc. Market makers can also extract information from the trades that are coming their way and from other market observations; as a result, they may possess information that is superior to that of the investors.

Market makers and investors are represented in the simulation as autonomous agents that behave according to their individual strategies: these strategies may be built in, or can arise as a result of learning or evolutionary selection. The built-in strategies were mainly derived from interviews with market experts (market makers in particular), as well as extracted from the market data. In the learning domain, we use neural networks and reinforcement learning to generate strategies for agents. This creative element is important because it allows us to investigate the possibilities resulting from strategies that have not yet been discovered by players in the real-world market.

## Results

Our results fall into a number of different areas. The primary results, reflecting the focus of our investigation, pertain to tick size effects, i.e., how the market's performance depends on the current tick size. However, we have also investigated a variety of other issues, and we present some of those results below.

### **Tick Size Effects on Price Discovery**

The model has produced some highly suggestive and unexpected results. Specifically, the simulation suggests that a reduction in the market's tick size (e.g., from \$1/16 to a penny) can reduce the market's ability to perform price discovery, particularly when parasitic strategies

such as "SOES bandits" and day traders are present in the market<sup>1</sup>. This is illustrated in Figures 1 – 6.




Figure 1. Fundamental value tracking in the presence of parasitic strategies (tick size \$0.01), yellow = offer price, red = bid price, green = fundamental value). For successful price discovery, the fundamental value falls between the bid and offer prices.

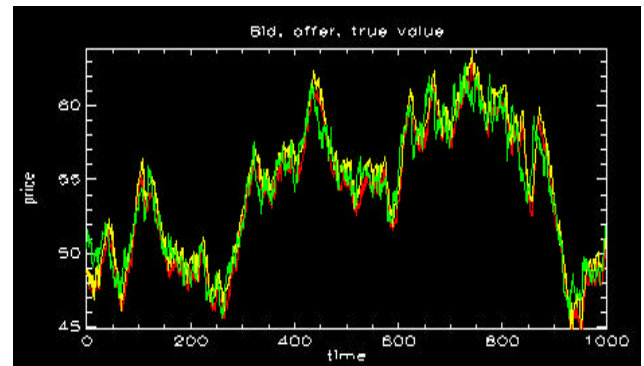


Figure 2. Fundamental value tracking with the same market conditions as in Figure 1, except that the tick size is \$1/16.

As Figures 1 and 2 demonstrate, in this particular scenario, market tracking is significantly improved when the tick size is increased from \$0.01 to \$1/16.

The analogous Figures 3 and 4 demonstrate that, under the same market conditions and with a smaller number of parasites, market tracking is almost unaffected by the same change in the tick size.

<sup>1</sup> A "parasitic" strategy in our framework is similar to the "wait in the background" strategy in [6]. It does not aim to provide liquidity to the market, but instead undercuts the current bid/ask price when the conditions are "right," and, after a transaction is completed, quickly withdraws from the inside market.

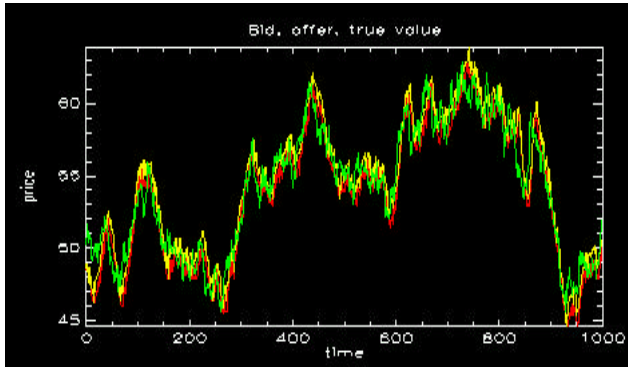


Figure 3. Fundamental value tracking with a few weakly parasitic strategies.




Figure 4. Fundamental value tracking with the same market conditions as in Figure 3, except that tick size is \$1/16.

To quantify those intuitive results, we have run similar scenarios 100,000 times. In particular, we calculated the standard deviation of the differences between the fundamental value and the average market price.

The results are presented in Figure 5. The value of the standard deviation of (price – fundamental value) for a population with an insignificant number of parasites remains basically flat. The same data, with a significant amount of parasitic market makers in the population, is different: up to a certain level, decreasing the tick size reduced the standard deviation of the price-value difference (effectively improving the market’s tracking and price discovery abilities); however, the standard deviation reaches a minimum at tick size \$1/32, and then starts rising sharply with further decreases in tick size.

This leads us to the conclusion that, in our simulation, a certain minimum value for the tick size is close to optimal in terms of the market’s ability to track the price. Decreasing the tick size below that level, in the presence of parasites, will greatly impair the market’s ability to perform price discovery.




Figure 5. The standard deviation of (Price – Fundamental Value), with a significant number of parasites (magenta) and a few weakly parasitic strategies (blue line).

In addition, we observed that most parasitic strategies also seem to create excess volatility in the market. These “excess volatility” phenomena are illustrated in Figure 6.




Figure 6. Market dynamics in the presence of a significant number of parasites. Tick size \$1/4.

We therefore conclude that there are two main effects of parasitic strategies on the price discovery process: (1) they impede the price discovery process, especially with smaller tick sizes, and (2) they seem to contribute to the creation of excess volatility in the market when the tick size is large. It is necessary to note that this excess volatility is present only in certain scenarios; it is affected by market conditions and the populations of investor and market maker types.

The above results were obtained for a population of investors who use only fundamentals to make their decisions. Thus there is always noisy information available, whether the market tracks the true value or not. However, ways to use this information to the agent’s advantage are not straightforward. Even if there are detectable deviations of the market price from the true value, they may be relatively short-lived, insufficient in magnitude, and also, most importantly, do not necessarily indicate the direction that the market price and true value will take in the next period. It may be possible to develop a strategy that

can take advantage of significant deviations of the market price from the true value, but in practice this does not seem to be easy to do and may not be theoretically always possible.

There thus appear to be two main impediments to the price discovery process in the model with fundamental investors:

1. **Insufficient information-processing capabilities of agents.** Adding learning dealers to the simulation somewhat improves market tracking, even in the presence of parasites, but not up to the point of eliminating any deviations from the true value. A more or less rigorous investigation of these phenomena is far beyond the scope of our research. However, it seems, as a plausible conjecture, that the same shortcomings may be present in real-world markets because the information-processing capabilities of a single individual are likely to be inferior to the information-producing capabilities of the market as a whole (which the individual is only a small part of). Another relevant phenomenon may be the insufficient reaction “speed” of hard-wired agents and insufficient speed of learning for the learning agents, which may also be present in the real world.
2. **Signal-to-noise ratio.** In our model, information is transmitted from the investors to the market makers through the price mechanism. Therefore, it is intuitively clear that relatively larger price moves will convey more information to the market. With a smaller tick size, however, the undercutting strategy that parasites normally use will produce less information to the market; in fact, small price moves may become undistinguishable from the noise. Potentially, effects of this kind may provide an upper limit for the efficiency of learning.

It would be interesting to see the effects of tick size on learning. If our conjecture about the interplay of signal and noise is correct, then tick size by itself may be able to alter the agents’ ability to learn and, through that mechanism, affect the market’s performance.

### ***Effects of Learning and Evolution***

For two reasons in particular, it is important to consider the results discussed above in the context of learning and evolution:

1. Learning and evolution should allow us to better calibrate the model so that it will be more representative of real-world markets.
2. Learning can produce types of behavior that are improvements over hard-wired agent behaviors in terms of realism and profitability. This is very important

for validating our model and for its potential applications.

### **Reinforcement Learning in Market Makers**

Many of the hard-coded dealer and investor strategies can evolve in the current system: when an agent goes bankrupt, the market selects one of the other agents and clones it, with more profitable agents more likely to be selected. A small random mutation is applied to the parameters of the new clone. If the mutated parameters lead to higher profitability for the new clone, then it is more likely to be selected for cloning in the future to replace a bankrupt agent. This constitutes a blind evolutionary search for parameters that result in high profitability. The search is blind because no feedback about performance is considered when adjusting parameters – the only feedback is from the agent’s survival and long-term profitability. If a mutation results in a good set of parameter settings, the agent will survive and be cloned, thus spreading its parameters; if the adjustment is bad, the agent will go bankrupt and die off, taking its parameters with it.

We have also incorporated reinforcement learning, a more sophisticated kind of learning involving immediate feedback, into our dealers. The reinforcement-learning dealers learn from their experience of different internal states and states of the market, and from the results of different actions. At each time step, the reinforcement-learning dealers are given a reward (or punishment) signal related to the change in value of the dealer’s total assets. In effect, these dealers learn the long-term profitability of different states or actions. After sufficient learning, the reinforcement learning dealers have some idea of which states and actions lead to greater rewards, and, at each time step, they choose the action that they think will lead to the greatest long-term rewards.

Dealers that can learn strategies are very important in the future development of the model because self-tuning strategies will increase market realism and the robustness of the agents under a variety of market conditions. Having adaptive agents in the model is also very useful for discovering unknown loopholes and unintended consequences of changes in market rules and agents’ incentive structures.

### **Types of Reinforcement-Learning Dealers**

We have developed several kinds of reinforcement-learning dealers. All share the same infrastructure; they differ mainly in what aspects of their internal state and the state of the market they pay attention to, and in what actions they can choose between. For these dealers, the possible actions are raising or lowering their bid and offer quotes by some amount. This creates a very large number of possible actions, so the actions are typically limited to raising or lowering quotes by a small amount

(e.g., one tick) or changing the spread of the quotes by a small amount.

Reinforcement-learning dealers take some time to learn, and they go bankrupt quite frequently at the beginning of learning. By default, they are resurrected after bankruptcy and are thus given the opportunity to continue learning, rather than being replaced by a clone of some other dealer. Each time a dealer is resurrected, this is said to be a new *incarnation* of the dealer. The most interesting region for the dealer to learn in is at the beginning of an incarnation, because there it is most vulnerable to mistakes. Once a dealer has become successful, it has more assets and is far less vulnerable to mistakes. Consequently, reinforcement-learning dealers can be configured to commit suicide if they survive for some number of market steps.

### The Reinforcement-Learning Framework and Incentive Structure

There are four major aspects to the reinforcement-learning framework for dealers in the model:

1. The reward structure – how short-term outcomes translate into reward or punishment.
2. The actions available to the dealer.
3. What information about the market and the dealer will be part of the state for learning.
4. The method used for learning what rewards to expect from particular states and actions.

The reward (incentive) structure is very important in determining the behavior of dealers in the stock market simulation. Different possible reward structures, and some of their consequences, are discussed below. For more information concerning the issues of and techniques used for Reinforcement Learning, refer to [8].

At each time step, the dealer receives a “reward signal.” The dealer attempts to learn actions that will maximize the value of the stream of reward signals expected in the future. However, the values of rewards in the future are discounted by an exponential factor based on the field “*discount factor*” in the agent.

The type of reward given to the agents can have a very strong effect on the strategies they develop. It is particularly difficult to create well-behaving dealer strategies, whose rewards are based on the value of their stock holdings at any given time; however, dealers whose reward signal is derived from an asset valuation based on the purchase price of stocks tend to be better behaved. This can be attributed to the stability of the reward signal, as the valuation of the dealer’s assets is insulated from market fluctuations and from the impact of the dealer’s own actions on the market value of the stock.

In order to evaluate the viability of the learning strategies with that of the hard-wired ones, we compared the assets of different market makers during the simulation. We found that, under a variety of conditions, learning strategies outperform the hard-wired ones. However, there are also market scenarios, in particular with many parasites or with high volatility, where the learning strategies at times find it difficult to survive.

### ***Fat Tails, Herding, Spread Clustering, and the Effects of Market Infrastructure***

Our model can help us understand how global characteristics of the market (such as interaction structure and rules) affect market behavior. These effects often propagate through the market by affecting individuals’ behavior, and only then does their effect appear on the global level. Interaction structure matters: for example, it seems that many effects we observe in the real-world stock market (such as spread clustering and fat tails) can be partly explained just by the market’s particular form of organization (existence of dealers, spreads, etc.).

It is clear that the market’s infrastructure and rules are among the most important determinants of global market performance. They are also important factors affecting the strategies of individual market participants, and the profitability of these strategies. As we noted above, often the outcome of a strategy is not uniquely associated with any particular feature of the model or the behavior of a set of market participants, but rather with a complex set of their interactions.

We have observed fat tails, persistent volatility, and spread clustering in our simulation under a variety of conditions. Interestingly enough, no explicit herding behavior is necessary for fat tails to appear in the simulation.

#### **Fat Tails**

Current stock market literature ([4], for example) asserts that variations in stock price follow a distribution that is similar to a Normal one but has a higher probability of extreme events (fat tails). In other words, fat tails are considered aberrations from normality in price data.

We investigated various scenarios with the underlying “fundamental value” following a log-normal walk. We have observed that, in many cases, the resulting market price dynamics resemble a Levy distribution rather than a log-normal one. For example, Fig. 7 is a histogram of the differences in logarithms of the fundamental value. This looks (and is designed to be) normally distributed<sup>2</sup>:

---

<sup>2</sup> The data for most of the histograms was derived from running the simulation for 100,000 periods. The details of most of the histograms (such as the fatness of their tails) will depend on the exact composition of the dealer and investor populations, as well as on other market characteristics.




Figure 7. Frequency histogram of the distribution of differences in the logarithms of true value

However, the analogous diagram of differences in the logarithms of the market price fluctuations observed in the simulation (Fig. 8) looks much more like a fat-tailed, or a Levy, distribution:




Figure 8. Frequency histogram of the distribution of the differences of the logarithms of the market price

A very interesting thing about the Levy pattern in Figure 8 is that it appears to be so orderly. Even the histogram of the underlying price dynamics has some glitches in it, but the histogram in Figure 8 appears to be extremely symmetric.

### Range of Fat Tails

We observed that fat tails tend to fade away with averaging the data, and even more so with increasing the time between the observations. For example, the following histogram is derived from the same data set as Fig. 8. The difference is that, in Fig. 9, the price data is sampled 50 periods apart.




Figure 9. Histogram of the differences in the logs of average price, with data points set 50 periods apart

Based only on visual analysis, it is very difficult to say whether the above histogram demonstrates deviations from Normal distribution at all.

## Conclusions

We have built a realistic model of a dealer-mediated market. We have focused our main effort on evaluation

of tick size effects. There are simple arguments that suggest that reducing the tick size may encourage competition and reduce the spread. However, there are more subtle issues concerning the effect of a reduced tick size on volatility, volume, the likelihood of crossed markets, the ease with which market makers can provide liquidity (in both small and large orders), the transparency of information flow in the market, and finally the efficiency of price discovery. Our results suggest that at least some of these “fuzzier” effects are adversely impacted by a reduction in tick size.

We have also observed in the simulation many important characteristics of real-world markets, such as the existence of fat tails, excessive volatility, and spread clustering.

We have also investigated learning market makers’ behavior and strategies, and found that, under a variety of conditions, artificial learning strategies outperform those extracted from the data or from the expert knowledge.

The model also allows investigation of a variety of other topics, such as the profitability of various market-making strategies in the simulated environment under changing market rules and conditions, the effects of ECNs and market fragmentation, etc.

## References

- [1] Cover, Thomas M., and Joy A. Thomas (1991): *Information Theory*, New York: John Wiley & Sons, Inc
- [2] Darley, Vincent M., and Stuart A. Kauffman (1997): “Natural Rationality” in *The Economy as an Evolving, Complex System* ed. W. Brian Arthur, David Lane, and Steven N. Durlauf, Reading, Massachusetts: Addison-Wesley.
- [3] Friedman, Daniel, and John Rust, eds. (1993): *The Double Auction Market: Institutions, Theories and Evidence*, Reading, Massachusetts: Addison-Wesley.
- [4] Lux, Thomas, and Michele Marchesi (1998): “Scaling and Criticality in a Stochastic Multi-Agent Model of a Financial Market,” University of Bonn Discussion Paper No. B-438.
- [5] Nasdaq (1998): *Market Making Guide*.
- [6] Rust, John, John Miller, and Richard Palmer (1993): “Behavior of Trading Automata in a computerized Double Auction Market”, in *The Double Auction Market* (above).
- [7] Smith Jeffrey W., James P. Selway III, and D. Timothy McCormic (1998): “The Nasdaq Stock Market: Historical Background And Current Operations,” NASD Working Paper 98-01.
- [8] Sutton, Richard S., and Andrew G. Barto (1998): *Reinforcement Learning: An Introduction*, MIT Press.