

Sixteenths or Pennies? Observations from a Simulation of the Nasdaq Stock Market^{*}

Vince Darley^a, Ph.D., Alexander Outkin^b, Ph.D., Tony Plate^c, Ph.D., Frank Gao^d, Ph.D.

Bios Group^{**}

January 07, 2000

^{*} 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.

^a Tel: 505-992-6737; e-mail: vince.darley@biosgroup.com

^b Corresponding author. Tel: 505-992-6743; e-mail: alexander.outkin@biosgroup.com.

^c Tel: 505-992-6712; e-mail: tony.plate@biosgroup.com

^d Tel: 505-992-6739; e-mail: frank.gao@biosgroup.com

^{**}
Bios Group, LP Phone 505-992-6700
317 Paseo de Peralta Fax 505-988-2229
Santa Fe, NM 87501 <http://www.biosgroup.com>

Introduction

We have built a model that represents a highly realistic picture (compared to the current state of the art) of a dealer-mediated market, with the flexibility to model many features of real-world markets. While we have conducted a fairly significant amount of research during the project, we have limited it to four areas:

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, at least in a stylized fashion, 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, exchange, investors, and dealers 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, but our main focus was 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).

In addition, the model has the flexibility to investigate various other questions — for instance, the effects of limit orders and market fragmentation, or developing better real-world dealer strategies — although some of these will require additional programming. The related issue of potential commercial applications of the model is of great interest to us. These applications are discussed separately at the end of the report.

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 — for instance, spread clustering, which occurs in our simulation.

The inherently complex nature of our model is extremely important to our results, and to the potential uses and applications of the model. It allows us to investigate the interplay between individual agents' features, their strategies, and the global characteristics of the market such as infrastructure, modes of information propagation, global constraints on

the interactions of market participants, etc. Some of the observed phenomena would be difficult if not impossible to discover using standard analytical techniques: for example, the fact that the patterns of spread clustering are affected by the investor population.

Model Overview

The main actors in the model are market makers (dealers) and investors, whose interactions within the marketplace produce price discovery and determine the market's dynamics.

The market contains a single security whose fundamental 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. Dealers 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.

Dealers and investors are represented in the simulation as autonomous agents, which behave according to their individual strategies: these may be built in, or can arise as a result of learning or evolutionary selection. 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.

Additional enrichments to the realism of the simulation include transaction costs, trading in a variety of volumes, trading in stocks with various capitalization, ECNs, differentiated information availability, various types of interaction between the players (such as preferencing), learning, and evolution. Some of these features are implemented in the present version of the simulation, while others are possible with additional development.

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 also present these 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. This is illustrated in Figures 1 — 6.




Fig 1. True value tracking in the presence of parasites. (Tick size \$0.01, yellow = offer price, red = bid price, green = fundamental price.) **In successful price discovery, the fundamental value falls between the bid and offer prices.**




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

As Figures 1 and 2 show, in this particular scenario, market tracking is significantly improved when the tick size is increased from \$0.01 to \$1/16. It is interesting to note that, in the previous versions of our simulation, a larger tick size increase was needed to demonstrate essentially the same effect. Surprisingly, the increased robustness of the market and the strategies made this effect stronger.

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.




Fig 3. True value tracking with a few weakly parasitic strategies.

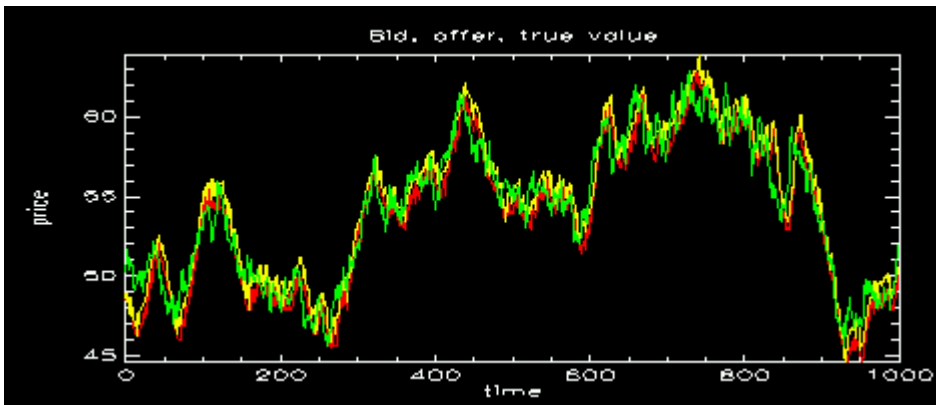


Fig 4. True value tracking with the same market conditions as in Figure 3, except that tick size is \$1/16.

In order to support these visually appealing pictures with hard numbers, we have run similar scenarios 100,000 times. The results appear in Table 1 of Appendix 1. In particular, we looked at the standard deviation of the differences between the true value and average market price. That data is presented in Figure 9.

As shown in Figure 5, the value of the standard deviation of (price — true 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 (thus effectively improving the market's tracking 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.




Fig 5. The standard deviation of (Price — True Value), with a significant number of parasites and a few weakly parasitic strategies.

One explanation for the increasing standard deviation of (price — true value) in a population with many parasites is the fact that the most parasitic strategies also seem to create excess volatility in the market. Visually, it looks as if the true value is roughly approximated by the dealer's quotes, but the swings become larger in magnitude and at times unnecessary from the point of view of the underlying true value dynamics. These phenomena are illustrated in Figure 6.

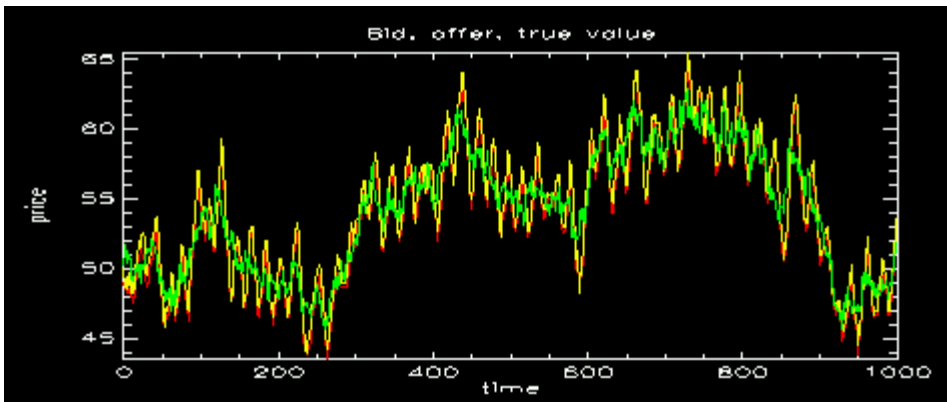


Fig 6. Market tracking in the presence of a significant number of parasites. Tick size \$1/4.

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 the 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.

One of the main factors affecting, and at times facilitating, the process of price discovery is the behavior of investors. Roughly, there are two main types of investors: those who make decisions according to the fundamentals, and those who do not. Of course, there are various shades of gray between these two, but the important factor is whether or not the investors rely on fundamentals to make their decision.

Throughout this paper we assume that investors have access to the fundamental value of the security, augmented by some noise. (We will also discuss the implications of various non-fundamental strategies on the market, but that will be necessarily more arbitrary for reasons of model calibration. Besides, adding non-fundamental types of traders tends in general to affect the price discovery process negatively, and thus strengthen our main results.)

These results were obtained for a population in which investors 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. Theoretically, we do not know if it is always possible.

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

1. 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. There are plausible reasons for this: (a) insufficient speed of learning, and (b) insufficient information-processing capabilities of the agents. A more or less rigorous investigation of these phenomena is far beyond the scope of our project. 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 is perhaps the insufficient reaction speed of hard-wired 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. Of course this is only a conjecture, based mainly on intuition and some previous experience; it seems, however, to be supported by the

data in Figures 7 and 8. We can see that, for a tick size of \$0.01, less detail of the underlying true value is reflected in the bid-ask spread as compared with a tick size of \$1/16. Potentially, effects of this kind may provide an upper limit for the efficiency of learning¹.

It is also interesting to note that the improved parasitic strategies not only seem to be more robust, but also, importantly, are much more profitable than before. However, we realize that in the real world there is an abundant supply of parasitic strategies that are not likely to succeed in the long run. Therefore, developing profitable parasitic strategies was not the main thrust of our investigation.

In addition, it is interesting that, even in the presence of large numbers of parasites, the market seems to be able to track the true value rather effectively. In Figure 3, we can attribute the relatively high standard deviation for large tick size values to the excess volatility produced by parasites (illustrated in Figure 6), and the deviation for small tick size values to the market's impaired ability to track the true value.

We have been able to greatly improve the market's ability to track true value by improving individual agents' abilities to perform price discovery. Nevertheless, there are scenarios in which non-tracking persists. This is especially true when there are relatively large populations of parasites and a relatively small tick size. Interestingly, with the improved population of traders in general, and parasites in particular, the adverse tick size effects become smaller (when measured as absolute deviation of true value from observed average price), but they are more pronounced with a small tick size. For example, in the current model the differences in tracking between tick sizes of \$0.01 and \$1/16 are much more apparent than they were in previous models.

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. We also may investigate further into whether the market is stable or unstable, in equilibrium or not; whether its evolution and learning tend to push it towards the edge of chaos; and whether a larger tick size may have stabilizing effects.

The tick size does affect the observations we have made about fat tails and normality on a larger time scale. Figures 1 through 6 were created with a tick size of \$0.01; the analogous histograms for a tick size of \$0.08 look similar, but, for example, the analog of Figure 5 appears less ordered and slightly skewed.

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 Investigating these types of questions is far beyond the scope and the spirit of the current project. It does, however, constitute a very intriguing and potentially commercially applicable research agenda.

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 Dealers

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 are presently incorporating 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. This signal is 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 the agents' incentive structures. There is still much work that can be done on investigating the behavior of learning agents and its effect on market behavior. One possible direction is to investigate the effect of the presence of learning agents on various aspects of market behavior such as price discovery, volatility, information flow, etc. Another direction is to investigate how learning agents respond to different incentive structures and different sets of market rules. This latter direction is potentially valuable for building simulation models that help us understand how real-life investors and dealers will respond to changes in market rules and to the new opportunities that are constantly being created in contemporary markets.

In the following section, we describe the different types of reinforcement-learning dealers we have experimented with, and some of our initial experiences with them. Subse-

quently, we describe the general framework of reinforcement learning and some of the issues that arise in using reinforcement learning in the simulation.

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. The reinforcement-learning dealer we have had the most success with is one that learns how to set its spread. We have also had some success with a learning dealer that learns how to move its quotes up and down, but we have not yet thoroughly investigated its behavior or its effect on the market.

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 *Spread Learner* is a version of the reinforcement-learning dealer. It is configured to learn how to set its spread between its bid and offer quotes in order to maximize profitability. It uses hardwired rules to adjust the center of its quotes to try to balance its inventory (similarly to the *Basic Inventory Dealer*), and learns how wide to place its quotes about the center. It always adjusts its spread by an even number of ticks in order not to interfere with the adjustment of the center of quotes. In its default configuration, the *Spread Learner* pays attention only to its spread (i.e., the spread is the only component of its state), and will thus try to find the spread that maximizes profitability under all conditions. With this simple configuration, it learns reasonably quickly; after several thousand time-steps, it begins to perform reasonably well. In more sophisticated configurations, it could also pay attention to other aspects of the market such as volatility or the buy/sell ratio; however, these more sophisticated configurations learn more slowly.

One interesting aspect of the *Spread Learner* is the *parity* of its spread: a spread that is an even number of ticks is said to have *even parity*, and a spread that is an odd number of ticks is said to have *odd parity*. Since the spread is always adjusted by an even number of ticks, the *Spread Learner* will maintain the same parity throughout an incarnation. The parity of the spread can be controlled with two options. The first is *Alternate spread parity* — if true, the *Spread Learner* will flip its spread parity between incarnations; if false, then the option *Even spread parity* can be set to true or false to give the *Spread Learner* an even or odd spread parity throughout an entire simulation run. In

some experiments an intriguing effect has been observed: *Spread Learners* with odd parity learn much more quickly than *Spread Learners* with even parity. We do not yet know how widespread this effect is, or the reason for it.

The *Q-Learner* is a more versatile type of reinforcement learning dealer. It learns the consequences of different actions in various states. In its default configuration, it pays attention to the following state aspects (at the level of detail specified in the parenthesis):
The number of buy trades (buys by investors) in the last time step (<2 or ≥2)
The number of sell trades (sells by investors) in the last time step (<2 or ≥2)
Its inventory (≤200, between —200 and 200, or ≥200)

The *Q-Learner* has ten different actions available to it:
Set its bid quote at or one tick below or above the current best bid, and likewise its offer quote (9 actions). Set its bid and offer at its maximum spread centered on the current best bid and offer

Because it considers many more actions and state aspects than the *Spread Learner*, this dealer takes much longer to learn. It learns to perform reasonably well after several tens of thousands of time steps.

The *Reinforcement Learner* is also a more versatile type of reinforcement learning dealer. It differs from the *Q-Learner* in that it learns the consequences of being in various states, and assumes that aspects of the state of the market that it cannot control (such as the quotes of other dealers, or the buy and sell volume) remain constant from one step to another. This assumption, though obviously false, allows the *Reinforcement Learner* to learn more quickly than the *Q-Learner*. In its default configuration, the *Reinforcement Learner* pays attention to the following state aspects:

The number of buy trades (buys by investors) in the last time step (<2 or ≥2)
The number of sell trades (sells by investors) in the last time step (<2 or ≥2)
Its own spread, in ticks
The position of its own bid relative to the best other bid (at, above, or below)
The position of its own offer relative to the best other offer (at, above, or below)

However, because the assumption of market constancy is false, the information learned is not as good as that of the *Q-Learner*, and in its default configuration the *Reinforcement Learner* does not perform as well as the *Q-Learner*. Note that the *Spread Learner* is actually a very simple specialization of the *Reinforcement Learner*: the only state component it pays attention to is the size of its own spread. Its success can partly be attributed to the fact that it knows exactly what effect its actions will have on its own spread.

The Reinforcement-Learning Framework and Incentive Structure

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

The reward structure — how short-term outcomes translate into reward or punishment.

The actions available to the dealer.

What information about the market and the dealer will be part of the state for learning.

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 described below. For more information concerning the issues of and techniques used for Reinforcement Learning, the reader is referred to Sutton and Barto's 1997 book.²

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. Suppose for the sake of this discussion that the discount factor is 0.9 (its default value is actually 0.95). The discounting means that the expected reward at the next time step is taken at full value, but the expected reward at the time step after that is discounted by 0.9, the expected reward at the time step after that one is discounted by 0.9^2 , etc. The sum of this stream of discounted future rewards is the expected discounted returns, which the agent tries to learn how to maximize.

The type of reward given to the agents can have a very strong effect on the strategies they develop: for some types of reward structure that seem reasonable at first, agents learn successful but quite undesirable strategies for maximizing the sum of discounted future rewards. The reward structures we have considered have two components: a method for calculating the value of a dealer's assets, and a method for translating the change in the value of these assets into a reward signal. The value of a dealer's assets is the sum of the dealer's cash and the value of its inventory. There are three methods for valuing stocks in the inventory held by a dealer:

The current market value, i.e., the mean of the best current bid and best current offer.

The current market value, but disregarding the dealer's own bids and offers.

The average purchase price of stock in the inventory, i.e., the price paid at the time stock was added to the inventory (with suitable modifications for negative inventory).

Once we have a method for calculating the value of the assets held by a dealer, we can derive a reward signal to be given at time step t , based on the assets at time step t (a_t) and the assets at the previous time step (a_{t-1}), in one of these ways:

The change in assets, i.e. $a_t - a_{t-1}$, or

The logarithm of the proportional decrease or increase in assets, i.e., $\log(a_t/a_{t-1})$

Both of these reward signals will be positive if the value of assets has increased, and negative if the value of assets has decreased. The second method is more appropriate when we expect that the earnings of a dealer should be proportional to its assets; the first is more appropriate when we expect the earnings of a dealer to be unrelated to its assets. In the current state of the simulator and agents, dealers have limited opportunities for trading in each time-step, so the absolute change in value of assets is more appropriate for determining the reward signal than the log change.

² References appear at the end of the paper.

The method chosen for valuing stocks can be critical. If stocks are valued at the overall bid/ask mean, including the dealer's own quotes, dealers tend to learn that the best course of action is to raise prices unconditionally. This results in the dealer's having a large positive inventory. The dealer achieves high rewards by constantly pushing up the value of the stock in its inventory. Eventually the dealer reaches such a highly leveraged position that it is very vulnerable to small movements in the market value of the stock, and the dealer goes bankrupt. At bankruptcy there is a large negative reward. However, the prospect of a large negative reward is not sufficient to discourage dealers from this strategy of price inflation because it is discounted. The same effect is also observed with price deflation: if a dealer has a large negative inventory, it can increase its asset valuation by deflating stock prices. However, this effect is observed less frequently, probably because there is a relatively close lower bound on the price of stocks. Dealers whose asset valuation is based on the best bid and offer quotes of other dealers still seem susceptible to attempting to inflate or deflate the price of stocks: even though their own quotes do not directly affect the valuation of their assets, they learn that, when they place higher quotes, this induces other dealers to also place higher quotes.

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.

A learning dealer tends to learn that a reward for using an even spread value can be much higher than a reward for using an odd one (or vice versa). Potential explanations of this phenomenon involve not only the individual dealer's behavior, but the aggregate behavior of the dealer population as well. For instance, if the population of dealers is using mostly even spreads, then it may be very detrimental for a new dealer to use odd ones (the market will move too fast against him, and inventory control will be very difficult). This is an interesting phenomenon that we can explore further for research as well as practical purposes — i.e., developing trading strategies that take this effect into account.

Fat Tails, Herding, Spread Clustering, and 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 does matter: for example, it seems that many effects we observe in the real-world stock market (such as spread clustering and fat tails) can be in part explained just by the market's particular form of organization (existence of dealers, spreads, etc.). Potentially, the infrastructure may be very important for the market's aggregate behavior, and this may have implications for other environments than markets.

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. Although these results will be presented in future papers, we give here a few examples of observations related to the fat tails.

Fat Tails

Current stock market literature (Lux and Marchesi, 1998, for example) asserts that variations in stock price follow a distribution that is similar to a Normal one but has, however, a higher probability of extreme events (fat tails). In other words, fat tails are considered aberrations from normality in price data. The precise meaning of normality is somewhat unclear — i.e., where is the assumed normality derived from?

In our simulation, normality is easy to derive — among other possible alternatives, we can investigate various scenarios with the underlying true price 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 true value. This looks (and is designed to be) normally distributed³:




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

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

³ 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 fatness of their tails) will depend on the exact composition of the dealer and investor populations, as well as on other market characteristics.




Fig. 8. Frequency histogram of the distribution of the differences of the logarithms of the average 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 absolutely symmetric!

Range of Fat Tails

We observed that fat tails tend to fade away with averaging the data (i.e., plotting a histogram of the moving average), 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.




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

The histogram in Fig. 9 seems to be much closer to a Normal distribution than the histogram in Fig. 8. In fact, 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. Clearly, difficult policy decisions must be made as to the correct balance between these different factors, and Bios simulation model sheds important insights that will help to make those decisions the correct ones, at least in the context of decimalization and changing the 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. Additional research is necessary to better understand those effects, in particular, calibrating the model (for example, market makers strategies) with real-world data (which is currently in progress).

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

Beltratti, Andrea et. al. (1996): *Neural Networks for Economic and Financial Modelling*. London: ITP.

Bigus, Joseph P., and Jennifer Bigus (1998): *Constructing Intelligent Agents with Java*. New York: John Wiley & Sons, Inc.

Christie, W.G., and P.H. Schulz (1994): "Why do NASDAQ market makers avoid add-eight quotes?," *Journal of Finance*, 49, 1841-60.

Cover, Thomas M., and Joy A. Thomas (1991): *Information Theory*, New York: John Wiley & Sons, Inc

Darley, Vincent (1999): Towards a theory of autonomous, optimizing agents. Dissertation, Harvard University.

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.

Friedfertig, Mark, and George West (1998): *The Electronic Day Trader*, New York: McGraw-Hill.

Friedman, Daniel, and John Rust, eds. (1993): *achuseons,s, F 1 T . TD T . nomy as an T Americd Lnancial MReview. 1*

Appendix 1

Table 1. Statistical Values Related to (Price — True Value)

Tick Size	Few Parasites			Many Parasites		
	Mean	Standard Deviation	Absolute Deviation	Mean	Standard Deviation	Absolute Deviation
100	0.0082	1.0831	0.7796	-0.062	1.425	0.8941
50	0.0204	1.0747	0.7711	-0.039	1.2127	0.8399
32	0.0103	1.0882	0.773	-0.025	1.1317	0.8227
25	0	1.0745	0.7682	-0.025	1.15	0.8491
16	-0.0012	1.0854	0.77	-0.024	1.1803	0.8978
10	-0.0078	1.0432	0.7522	-0.012	1.233	0.9579
8	0.0192	1.0837	0.7708	-0.024	1.2481	0.9696
4	0.0197	1.0713	0.7615	-0.015	1.3539	1.0598