

## Chapter 1

# Foresight, Unpredictability and Strategies

### 1.1 Introduction

The difficulties of understanding financial markets have fascinated scientists, market practitioners, and laymen alike as long as the markets themselves have existed. In addition to well-known theoretical considerations, the main impediments to understanding the markets' behavior include the sheer complexity of the market infrastructure, rules, and behaviors of the participants.

The focus of our approach is to understand those components by modeling them individually and by creating a simulation of the market where those individual components interact with each other and the resulting market behavior emerges through accumulation of those interactions over time. We employ agent-based modeling to achieve this purpose.

We attribute the germination of this approach to two important recent scientific and technological breakthroughs:

- (1) Changes in conceptual paradigm - due to development of non-linear dynamics, chaos theory, game theory, agent-based approaches, the world is more and more recognized as a complex adaptive system, which dynamics arises due to many factors and their interactions - agents, their behaviors and goals, interaction rules, etc. Such systems exhibit emergent properties - patterns of dynamics behavior that arise from interaction of their components, which (patterns) by themselves are not properties of those components<sup>1</sup>. Behavior of such systems may not be well approximated by equilibrium conditions often assumed in economics or finance; instead the interesting and practically important dynamics may lay out-

---

<sup>1</sup>See Mainzer (2004), Nicolis and Prigogine (1989), Kauffman (1995), and Haken (1988) for discussion and examples of complex and complex adaptive systems.

side the equilibria regions, or it may be altogether chaotic. Perfectly rational representative agents approaches so adored by modern economics are insufficient to describe such a system. Instead, a butterfly flapping its wings in the Amazon causing a hurricane in the Caribbean is now a recognizable household phenomenon. In this view, the world is a decentralized parallel system, where many actions take place at the same time, often autonomously or independently. Decisions are made and actions are taken by heterogeneous agents, interacting with each other in spatially localized, time-dependent contexts. Complete information is rarely available, instead it is often local, incomplete, and noisy<sup>2</sup>. This is the world where actions of many participants, rather than of a single Walrasian auctioneer determine the market clearing prices. Yet, this newfound paradigm did not bring sufficient analytical tools to analyze this degree of complexity, and computer simulations become necessary.

- (2) Relatively recent developments in the computer hardware and software just recently reached a point where useful simulations of systems described above become possible. It is a confluence of many advances: computers' raw speed and memory, object-oriented technologies, development of software frameworks and libraries, including agent-based frameworks, parallel processing, wider super-computer availability, advances in software design, and many others.

The end result of those developments is that practically useful modeling and analysis of complex systems has become possible, albeit difficult.

The unique feature of our approach is that we explicitly model the behaviors and the decision-making process of individual market participants such as market makers and investors, interactions between them, and the market rules and infrastructure. The simulation model allows investigation of market behaviors under a variety of scenarios and conditions. The goal of our approach is not only do describe what have happened, but also to model how (through which interactions, strategies, rules) did it happen, and attempt to understand why did it happen.

As a result we do not need to use the market equilibrium assumption so often relied upon in modern economics. Instead, market dynamics arise

---

<sup>2</sup>Albin (1998) stressed the importance of using appropriate tools for modeling information: "The core of my argument is a claim that informational technologies and processes have critical properties which are imperfectly captured by models in the mechanistic analytical tradition but which are well defined when building blocks of the economy are represented as abstract automata."

as a result of interaction of the market's individual components. Nevertheless, we show in Chapter 2 how our approach can be used to replicate the Glosten-Milgrom model, where market dynamics converges to an analytically tractable equilibrium.

Our approach allows us to investigate scenarios where the market dynamics may deviate from equilibrium and situations where the market changes suddenly, or undergoes a phase transition, or other dramatic, qualitative change.

This 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. The ability to investigate disparate agent strategies is particularly important – as we will see later, strategies employed by markets participants are one of the most important factors responsible for the particular dynamics of the market. We define a strategy of a market participant as a set of rules, governing that participant's behavior. An example strategy could be as simple as “buy if the price is going up” and “sell if the price is going down”. More sophisticated strategies are discussed later in the book. Some of the observed phenomena would be difficult if at all possible to discover using standard analytical techniques: for example, the fact that the patterns of spread clustering are affected by the investor population (see more on that in the chapter on spread clustering).

As our starting point of departure, we investigated the possible effects and outcomes 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. On calibrating the model, we discovered that it exhibits a number of behaviors normally associated with real-world markets, 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 expert knowledge. This creative element is important because market regulators are especially interested in strategies that have not yet been discovered by players in the real market, motivated by their goal of designing a regulatory structure with as few loopholes as possible, in order to prevent abuses by devious players.

We also found, by analyzing the real world data, that the strategies, employed by the Nasdaq market makers underwent a phase transition as a result of the previous tick size change. Namely, when the tick size was changed from  $\$1/8$  to  $\$1/16$  in 1997, the strategies employed by the market makers underwent a quantitative and qualitative change.

We observed this change by representing strategies of all the market makers as a function of strategies of two representative market makers. That means that the strategy of every market maker in the market can be expressed as a linear function of two parameters, each expressing dependency on one of the representative market making strategies. Investigating how those parameters change with tick size allows understanding whether a phase transition in market makers' behavior has happened concurrently with the tick size change.

There are two ways for this parameter change to occur:

- The coefficients (before and after) significantly overlap in the space. Here there are no "phase transition"-like changes in the market; the transition is smooth. This type of behavior (transition) could be predicted with good accuracy.
- The coefficients do not overlap, or the non-overlap area is significantly large. In this case we have an abrupt transition, like the phase transitions in physics. Now prediction becomes problematic.

As will be elaborated upon further in this book, there is strong evidence that a phase transition in the market maker strategies occurred at the same time as the tick size changed in May-June 1997.

If the phase transition result is generic, then another change in tick size will lead to a phase transition in strategies. Therefore, the population of strategies before and after the tick size change must be quantitatively and qualitatively different.

This is on one hand an intuitive, and on the other hand mind-boggling result: some, even a-priori inconsequential changes to market rules, such as changing the tick size from  $\$1/8$  to  $\$1/16$  may produce market changes that would be impossible to predict beforehand. What is more, the end result, being the outcome of a phase transition, may depend not only on the change itself, but also on minor details on how the change would be implemented and carried through. Therefore, the results of those changes can be unpredictable.

This result establishes serious limits on how much prediction of the effects of a particular change can be done before the change happens. Yet,

we believe our approach allows investigating what could happen given the strategies of the market participants employed at present, given the ability to create a new population of strategies using learning and evolution, and given the rules and interactions in the market place (and changes to them) as they are understood by the researchers and implemented in a model. While such approach may not provide a definitive prediction this is not a problem for two reasons: 1) As outlined above, a definitive prediction may not be possible for example in case of phase transitions; 2) This approach allows investigating a range of “what-if” scenarios, thus providing a capability to plan and adapt.

These considerations do not however, invalidate the prediction in general – rather, they impose limits on what can be predicted. Even if the details of the evolution of a population of strategies may not be predictable, certain system-level properties may be predictable. We believe, the majority of the predictions we’ve made using our model fall into the latter category (see more below, and in Chapter 8).

Perhaps as important, predicting social systems in general may not be a self-consistent exercise - once a prediction is made and known, it is likely to affect the system behavior. Such changes may invalidate the prediction, make it self-fulfilling, or generally change the interactions or perceptions of system participants, so that the original model may become invalid.

Additionally, a certain degree of unpredictability is necessary for the market to operate – after all a buyer and a seller to a stock transaction are unlikely to have the same views of the stock being traded. A need for uncertainty, as a prerequisite for a viable market or a society in general is expressed in Peters (1999). Additionally, there is an interesting interplay between uncertainty, as connected to entropy, and information generation - for example the real-world ant algorithms implemented by the real-world ants work in part because the information left in a particular location in the form of a pheromone trail dissipates to other locations - and thus can be discovered more easily (see Parunak and Brueckner (2001) for example).

Interestingly, extra randomness may improve behaviors of individual trading strategies in our simulation. For example, the ‘Volume’ dealer, described later, earns better profit when a degree of randomness is added to its decisions. Perhaps this is analogous to simulated annealing algorithms that require a degree of randomness to avoid getting stuck on a suboptimal solution.

While building the simulated model of the market, we interacted extensively with many market participants: market makers, brokers, traders,

large investors, etc. We found this interaction invaluable - as a source of information in particular on often subtle details of market operations, as a venue for for verifying our assumptions and simulation results, and at times as a source of constructive criticism. One conversation with a market maker still stays clear in our minds. He was supportive, but skeptical. The core of his skepticism lay in this question: how one can model the fear and greed often ruling the market behavior? This is a valid point: while fear and greed affect markets immensely, as has been time and again demonstrated by numerous booms and busts, understanding of underlying individual and mass psychology is lacking.

In our approach we address this problem by explicitly modeling strategies of individual market participants, by allowing those strategies to evolve over time due to individual learning or evolutionary selection, and by allowing to investigate various what-if scenarios by using user-defined strategies.

In the user-guided simulation, Market Maker and investor agents (institutional investors, pension funds, day traders, casual investors) buy and sell shares using various strategies. The agents' access to price and volume information approximates that in the real-world market, and their behaviors range from very simple to complicated learning and evolutionary strategies. The simulation can be managed and run with a user-friendly graphical interface, or it can be run in batch mode to produce larger amounts of data more quickly. All results are easily graphed and compared for effectiveness.

This approach has proven to be successful: many of our predictions about the market's behavior after decimalization came true.

The following list summarizes our most important predictions:

- (1) Decimalization (tick size reduction) will negatively impact the price discovery process.
- (2) Volume will increase, potentially ranging from 15% to 600%.
- (3) Ambiguous investor wealth effects may be observed. (Investors' average wealth may actually decrease in the simulation, but the effect is not statistically significant).
- (4) Phase transitions will occur in the space of market-maker strategies. (This was observed in the transactional data).
- (5) Spread clustering may be more frequent with tick size reductions.
- (6) Parasitic strategies<sup>3</sup> may become more effective as a result of tick size reductions.

---

<sup>3</sup>Parasitic strategies are defined and discussed later. Briefly, we say that a strategy is parasitic when it attempts to use information and liquidity provided by other market

Overall, our predictions are supported by available studies of decimalization results (see, for example, Nasdaq Economic Research (2001) and Chapter 8):

- There is strong, albeit indirect, evidence that reducing the tick size negatively impacts the price discovery process.
- There are strong indications that not all investor types benefited from decimalization. (There are ambiguous wealth effects, like those observed in our simulation.)
- There is unambiguous support for our prediction of phase transition in market maker strategies as a result of decimalization.
- Nasdaq Economic Research (2001) reports that there is a significant natural quote clustering in the decimalized market.
- That same source reports that there is evidence of more stepping in ahead of the customers' orders (in particular on electronic communication networks, or ECNs) and more time when ECNs are on the inside market, which we interpret as indications of increased activity of parasitic strategies.

Our volume increase prediction had a wide range of possible outcomes (an increase from 15 % to 600%). The upper range of that prediction has not happened immediately after the decimalization. However, one needs to take into account that the decimalization occurred in a down market, where investors had little appetite for increased trading.

The outcomes of decimalization not only confirmed our predictions about the behavior of the aggregate market, but also confirmed the validity of our models of investor and market maker behavior. In particular, we predicted that parasitic strategies may become more prominent after tick size is reduced and that the spread, or quote clustering patterns are likely to appear more often. Both of those predictions are strongly supported by the available data.

In addition, anecdotal support for our conclusions has come from Wall Street practitioners who have read some of the research reports from which this book arose. As a general approach, market makers on Nasdaq have shifted their behavior from an "old world" in which their primary role and activity was a positive one of providing liquidity to the market (and being

---

participants, without providing those itself. We do not offer a value judgement by calling a strategy "parasitic", but rather use commonly accepted terminology.

compensated for that role and risk by the spread), to a “new world” in which their primary activity is more negative, usually extracting liquidity from the market by front-running and similar strategies.

The next step is to bring this type of modeling from the realm of research to the practice of the real-world markets. We believe models of this kind can be useful for both regulatory and trading or market making purposes because they offer unparalleled ability to understand market dynamics, incorporate user knowledge and proprietary data sources into the simulation and to calibrate the model in the real-time.

## 1.2 What is Agent-Based Modeling?

Here we present a brief summary on how we define and use agent-based modeling. Additional sources include Epstein and Axtell (1996), Schelling (1978), Shoham (1993). See also section 1.5 of this Chapter.

In agent-based modeling, appropriate parts of the complex system are modeled as autonomous decision-making entities, called agents. Agent-based modeling does not require representing everything as agents, only the parts of the system that may benefit from it from the modeling perspective. At the simplest level, an agent-based model consists of a system of agents and the relationships between them. Each agent individually assesses its situation and makes decisions based upon a set of rules. Agents may execute various behaviors appropriate for the system they are a part of - for example, producing, consuming, or selling. Repetitive, competitive interactions between agents are a feature of agent-based modeling, which relies on the power of computers to explore dynamics out of the reach of pure mathematical methods.

Agent-based modeling is a simulation methodology which, when carried out carefully, results in a realistic simulation of a system because it emulates the manner in which the world really operates, instead of using such assumptions as equilibrium or perfect rationality. Even a simple model can exhibit complex behavior patterns and provide valuable information about the dynamics of the real-world system that it emulates. This emulation can include real world issues such as limited information, noisy data, limited computation, rules of thumb and so on. In fact it would be difficult if not impossible to create a non-trivial agent-based model with agents exhibiting the standard *homo economicus* rational expectations rules with perfect rationality, infinite computation, etc.

In addition, agents are capable of evolving, allowing a different set of unanticipated behaviors to emerge. Agent-based modeling can incorporate recent advances in neural networks, machine learning, and genetic algorithms to allow realistic learning and adaptation. Such models can be used in decision-making, as well as to develop flexible, adaptive strategies in dynamic business environments.

An agent-based simulation is a discrete time simulation, where agents make decisions at certain time steps, taking into account actions of themselves and of other agents, as well as other external and internal factors. Real-world entities that are important from the modeling point of view are often modeled as separate software objects.

We call an object a self-contained programming entity that characterized by its state, its interaction with other objects and information it can receive, process and transmit (see Roy and Haridi (2004) for a precise definition and extensive discussion of objects and other programming concepts).

We call an agent an autonomous entity that has goals or objectives, ability of receiving information, and making and executing decisions while working towards those objectives. Often this decision-making ability is called agent strategy, and in a simple case can be represented as a set of rules of what action to take under what conditions. An agent attempts achieving his goals by interacting with other agents. He receives signals from other agents and external sources and makes his own decisions on what actions to take. Agents are often programmed as separate objects or collections of objects.

We normally model as agents the entities that have relatively complicated decision-making rules (behaviors) or represent most interest from the modeling point of view. External factors are often modeled as individual objects, but not necessarily as agents<sup>4</sup>.

Generally, an agent-based simulation contains the components listed below. This list is intended mainly as an illustration of a set of concepts useful when implementing an agent-based model, and not as a mathematically precise definition, therefore, some of the categories below may not be exclusive of each other. Furthermore, the same terms may be defined differently in different sources. Review of those differences is outside of the scope of this book. More theoretical approaches to agent-based modeling exist as well. For example Shoham (1993) describes agents as having

---

<sup>4</sup>At this point it is clear that agent-based modeling maps naturally into object-oriented programming. However, agents can also be implemented using other programming techniques as well.

the following attributes: belief, decision, capability, obligation. We find the concept of obligation particularly useful for modeling various financial contracts that span non-trivial periods of time - short sales, futures and options, repo markets contracts, etc. For example selling short incurs an obligation to deliver security back within a certain period of time.

(1) Agents

- (a) Set of actions available (buy, sell, change contract price, etc.).
- (b) Decision-making rules and algorithms (how agent makes decisions; say, sell when the price of stock starts to go down and buy when the price starts to go up). Those decision-making rules are often called agent strategies. In practice, they can be implemented as separate strategy objects, and used by agents by means of strategy pattern (see Gamma et al. (1995) for example).
- (c) Scheduling of decisions and actions - at which intervals decisions are made (say, every week), or depending on which signal an action may be executed (say after a trade occurred). Different agents may have different time scales - some can make decisions every day and others every month or every second. An agent's schedule may incorporate actions the agent may choose to take (say analyze last hour's prices and volumes) as well as obligations an agent may have to fulfill (say return the shares he borrowed). Scheduling, however simple it may sound, is a critical component in agent-based simulations. It affects such aspects of the simulation as repeatability of results, determinism, and effects of randomness, particularly when simulation uses a distributed architecture.
- (d) State - how agent's state is described - amount of capital, inventory, production capabilities at present, etc. Most often the agent's state is changed as a result of interaction with other agents or with the environment.
- (e) Information in- and out-flows - what information agents take into account when making decisions, what information they provide to others.
- (f) What other agents/external factors this agent interacts with or affected by.

(2) Interactions and information flows

- (a) Who is interacting with whom.
- (b) What are the flows of information.

- (c) What are the results of agents' interactions: monetary payoffs, the commodity, money, etc flows. Example: say, agent A interacts with agent B, provides information to agent C and pays money to agent B.
- (3) Data available (what historical and other data is available for the whole system and individual agents).
- (4) External factors (environment). Normally those are the factors over which agents do not have direct control, but may want or may have to take into account when making their decisions. For an individual investor agent such factors can be the market price of security, Federal Reserve board interest rate changes, etc. - all those factors that an agent is unlikely to be able to affect through his actions and therefore may have to take as given.

Note that the very nature of modeling interactions between the agent participants in the system is a step away from most classical economics. The ability to model interactions explicitly allows the system to contain or exhibit localized, time-dependent events, and to understand how those events might propagate throughout the system.

For many practical and theoretical purposes we consider agent-based modeling an extension of game theory<sup>5</sup>. Similarities include: concepts of agents, strategies, payoffs, information flows, and interaction structures. The main differences are as following:

- The main goal of the agent-based modeling is understanding the dynamics of the system, while investigation of various equilibrium concepts appears to be the main goal of game theory.
- In a game-theoretic setting, agent payoffs are explicitly specified by a payoff matrix that is a function of strategies played by a population of agents, while in an agent-based model agents' payoffs usually arise as a result of interactions between agents, and often are not represented by an explicit payoff matrix.
- Game theory strives to quantify equilibria states of games analytically. Agent-based modeling usually uses simulation in order to understand dynamics of the system.
- Historically agent-based modeling has been applied to a broader set of

---

<sup>5</sup>Game theory is often called a theory of strategic interaction. It was introduced in von Neumann and Morgenstern (1944). An overview of game theory is outside of the scope of this book. See, for example, Fudenberg and Tirole (1991) for an introduction.

domains and to a broader set of problems than game theory.

These similarities can often be used productively: for example creating a simplified version of an agent-based model that can be analyzed using game-theoretic tools can be useful in understanding and validating the behavior of the more complex agent-based model.

### 1.3 Nasdaq Market Overview

Below is a brief overview of the Nasdaq stock market. We do not attempt a comprehensive description of the market here, but instead aim to provide a parsimonious description of the most essential features of the market to help the reader with understanding the simulation framework described in the book. For a more comprehensive overview, see Harris (2003) and Smith et al. (1998), among others. See also Friedfertig and West (1998) for a practical discussion of trading systems and interplay between trading strategies, trading systems, and market rules. O'Hara (1995) provides a classical overview how markets, such as Nasdaq, work from the theoretical point of view.

The Nasdaq stock market has undergone a period of rapid change over the last five or tens years, including decimalization, the rise of ECNs, significant rule changes, acquisitions, and competitive moves – many of them having significant implications on market dynamics, the market's inner workings, and trading. The Nasdaq Market's own websites, including [www.nasdaqtrader.com](http://www.nasdaqtrader.com) provide a significant amount of information on those recent changes.

Nasdaq is a dealer-mediated, quote driven market. On quote-driven markets, market makers post quotes (prices and volumes at which they are willing to conduct a transaction) and the investors may decide to accept those quotes. By posting those quotes, market makers provide liquidity to the market by allowing investors to execute transactions where the market maker is a counter-party, and by taking associated risks at the same time. Liquidity provision is also market-makers' fiduciary duty according to the Nasdaq rules and regulations. Other market participants include investors and brokers. Investors trade for their own account, and brokers can execute orders on behalf of other parties, such as investors.

Prices, at which the market makers are willing to buy or sell the securities to the public are called bids and asks respectively. The highest market

maker buy price and the lowest sell price are called the best bid and best ask respectively. The difference between the best ask and the best bid is called bid-ask spread, or just spread. Spread is an important measure, because it affects the profits made by the dealers and the costs incurred by the trading public<sup>6</sup>.

Tick size is the smallest price increment. Historically, on Nasdaq and other US markets, tick sizes were quoted in fractions (for example  $\$1/8$ ). Until 1997, the Nasdaq market tick size was  $\$1/8$ . In 1997 it was changed to  $\$1/16$ . Our research started in 1998 with the goal of understanding how future tick size reductions may affect the markets. Decimalization (a change to expressing quotes and prices in decimal system, rather than in fractions) and a tick size reduction were implemented by Nasdaq in Spring of 2001. Decimalization affected, directly or indirectly, the market dramatically. Some of its effects are discussed in Chapter 8.

Nasdaq also has order-driven components – for example ECNs and in some cases the market makers display public (individual investors’) limit orders. Therefore, there are two types of orders – market orders and limit orders. Market orders require execution at the best available price, limit orders, on the other hand have either the maximum price the order originator is willing to pay in case of a buy order or a minimum price the order originator is willing to accept in case of a sell order (thus “limit”).

In addition to being regulated by government agencies, such as the SEC, Nasdaq is also regulated by the NASD (National Association of Securities Dealers). When we began our research NASD was a parent organization for Nasdaq, and Nasdaq was a self-regulating entity. Back in 1996 NASD embarked on radical changes that would allow NASD to concentrate on market regulation and investor protection, and for Nasdaq to become responsible for market operations. At present this process is complete, and Nasdaq operates as a separate entity. More information can be found on NASD and Nasdaq websites and elsewhere.

## 1.4 Nasdaq Simulation Model Overview

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

---

<sup>6</sup>After decimalization and tick size reduction were implemented by Nasdaq in 2001, dealer profits earned from the spread (and the spread itself) generally declined significantly.

model many features of real-world markets. The main actors in the model are dealers, whom we call interchangeably 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 sometimes referred to as "true value", is exogenously specified: the underlying value of the security is assumed to fluctuate according to a random process. We will use "fundamental" and "true" value interchangeably throughout this book. In the context of the simulation, "true" value is perhaps more appropriate - while in the real world fundamental value is not directly observable, its analog is observable in the simulation - thus "true" value.

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 investors 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 in the simulation do not have access to the fundamental value of security, but can extract information from the trades that are coming their way and from other market observations; as a result, they may and most of the time do possess information that is superior to that of the individual 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.

In our research we initially 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.

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.

Through deeper analysis of systems like this, and a commensurate greater understanding of the relationships between liquidity, volume, volatility, etc, it may well be possible to build simpler models than ours which bring the results down to a more fundamental set of criteria. This is a promising avenue of current research, with Doyne Farmer and his collaborators perhaps bringing the best combination of practical experience and theoretical expertise together in promising new ways (also see the section 1.5 of this Chapter).

## 1.5 Existing Market Modeling and Simulation Work

A great many researchers in economics, finance, and related areas have investigated different aspects of financial markets or markets in general. Our approach is inherently multidisciplinary. Areas of inquiry relevant to our approach include not only financial markets, but also agent-based modeling, game-theory, machine learning, simulation design and analysis, software architecture, visualization, and others. We will not attempt a

comprehensive overview of relevant literature here, but rather point to a few sources that we found useful to our research while we were conducting it or after the fact<sup>7</sup>. Our main goal here is to outline the areas of thought, confluence of which contributed to the germination of our approach.

### 1.5.1 *Complexity and Agent-Based Modeling*

The agent-based literature has grown exponentially over the last two decades. We by no means attempt a comprehensive overview here, but instead list works that affected our approach to the Nasdaq stock market modeling. Epstein and Axtell (1996) outlines a cohesive framework for applying the agent-based modeling to social sciences. An earlier work by Schelling (1978) demonstrated the importance of understanding decentralized interactions, and their significance in determining the dynamics of the system. Albin (1998) demonstrates an approach to economics problems where complexity of interactions and agents ability to formulate and execute strategies are explicitly taken into account. Mainzer (2004), Nicolis and Prigogine (1989), Kauffman (1995), and Haken (1988) provide a discussion of various theoretical and practical aspects of complex systems. Shubik (1999) applies game theory to multiperiod models of economic exchange. A more light-hearted account of complexity in application to economics and accounting can be found in Brown and Palmrose (2005). Weiss (2001) provides a comprehensive overview of various issues in distributed artificial intelligence.

Our own Ph.D. theses provided an initial impetus for our approach to the Nasdaq stock market modeling: Darley (1999) considers a population of interacting agent strategies, adapting over time, selecting the most predictive strategy for each agent's environment, where the complexity of strategy selected is itself selectable. In such a situation an agent has the potential to evolve towards more and more sophisticated strategies (in some sense approaching the perfectly rational homo economicus). It turns out that, under quite general circumstances, it is not beneficial to adopt such complex strategies, and a simpler, boundedly rational set of agents perform better.

Outkin (1998) dissertation goal is to expand the game theory into domain of evolving, potentially non-equilibrium systems. He argues that to

---

<sup>7</sup>Most of the research described in this book was conducted during 1998 - 2001. Particularly earlier in that time period not much literature on agent-based modeling was available.

achieve that goal, adaptive strategies (defined as strategies that use rational as well as heuristic rules and change them according to the player's experiences in the game) and local interactions have to be modeled. He analyzes effects of local interactions and imitation strategies in the Prisoners' Dilemma game, and finds out that in an equilibrium where cooperation co-exists with non-cooperation (mixed equilibrium), there are cooperators who receive better payoffs than any non-cooperator. Using simulation studies, it is shown that dynamics of such models can exhibit chaotic behaviors, self-replicating structures ("comets"), as well as convergence to both complete cooperation and non-cooperation (see also Outkin (2003)). The concluding part of the dissertation analyzes analytically the stochastic dynamics of a population of locally interacting best response players and shows that stationary distribution is of the Gibbs-Boltzmann type and the long-run equilibria are the maxima of a potential function.

### 1.5.2 *Financial Markets and Agent-Based Modeling*

Topics of relevance to our project have included diffusion of information in markets (Kyle, 1985; Grossman and Stiglitz, 1980), efficiency and statistical modeling of markets (Foley, 1994), speculation and bubble formation (Tirole, 1982), and many others.

There is a number of sources that provide a thorough introduction to how certain markets work in reality. For example Harris (2003) provides an extensive overview and a depth of detail on various markets including Nasdaq, participant strategies, and rules governing those markets. Another example, mainly concentrating on the FX markets is Dacorogna et al. (2001). Friedfertig and West (1998) provide a useful introduction to the Nasdaq stock market from the point of view of day traders. We also found this book useful for modeling the day trading strategies.

The Market Making Guide (Nasdaq, 1998) provides useful information on how the Nasdaq market makers operate. Smith et al. (1998) provides an overview of the Nasdaq history, and the way it was operating in the late 90th. However, Nasdaq changed dramatically since then, in part because of the decimalization, but also because of competitive pressures, conversion to a publicly traded entity and concomitant split into Nasdaq Inc. (publicly traded organization), and NASD (a regulator). Current information, including market rules and regulations can be found on the Nasdaq and NASD websites ([www.nasdaq.com](http://www.nasdaq.com), [www.nasdaqtrader.com](http://www.nasdaqtrader.com), [www.nasd.com](http://www.nasd.com)).

Statistical approaches to markets have become particularly prevalent with Econophysics. A good introduction to statistical approaches in finance can also be found in Bouchaud and Potters (2000). Farmer et al. (2005a) demonstrates power of stochastic zero-intelligence approaches for understanding market behavior. Another example of zero intelligence approach can be found in Farmer et al. (2005b).

This study was initially motivated by existing work in market microstructure theory (O'Hara, 1995). This body of work on its own is quite compelling, yet it suffers from a lack of means for experimental verification (with some important exceptions, such as Rust et al. (1993)). For example, in an earlier version of our model, we investigated the behavior of the Glosten-Milgrom (GM) model, which examines the relationship between dealer spreads and the proportion of informed investors in the market. While this model is intuitively quite reasonable, it is very difficult to quantify the fundamental concept of "proportion of informed investors in the market." Ideally, one would prefer a numerical and experimental framework in which one could examine the regions of validity of analytic models like the GM.

Complexity and agent-based related approaches to the markets include Lux and Marchesi (1999, 2000); Chen et al. (2001); Bak et al. (1996); LeBaron; Levy et al. (1995, 1994); Caldarelli et al. (1997). Rust et al. (1993) outlines a model where various strategies interact in a market with a centralized auctioneer. They show that parasitic strategies are often the most successful. Axtell (2005) provides theoretical foundations for this decentralized view of the markets, where prices and market dynamics arise as a result of decentralized interactions between agents, rather than due to a Walrasian auctioneer. He also shows that Walrasian type of market clearing is not even feasible due to the non-polynomial complexity of algorithms. Axtell (2003) outlines problems with traditional economics approaches and argues that society can be represented as a large-scale computational system, that can be modeled using agent-based approaches. Das (2003) describes a market making algorithm based on "probability density estimate of the underlying value of a stock". Farmer and Joshi (2002) discusses the price impacts of various trading strategies in a model with market maker price formation. See also Darley, Outkin et. al. (2000, 2001) for our earlier descriptions of this research. Campbell et al. (1997) provides an excellent overview of more standard econometrics approaches to finance.

### 1.5.3 *Artificial Intelligence, Machine Learning, and Other Approaches to Modeling Agent Strategies*

Sutton and Barto (1998) outline a machine learning framework similar to that we've used to enable our market makers to learn. Bigus et al. (2001) provide an extensive summary of various techniques useful in agent-based modeling together with implementation examples. Beltratti (1996) provides examples of using neural networks in modeling financial markets.

At present there is a large number of various agent-based frameworks with varying functionality, addressing various design and performance aspects of agent-based modeling, many of which are suitable for using in building market models. Swarm was perhaps the first general agent-based modeling framework, developed at the Santa Fe Institute in the early 1990s. This was followed by other frameworks such as Ascape, which was created at Brookings Institute in the 90s, RePast, created at the University of Chicago, MASON, created at the George Mason University, and others.