TUTORIALS IN
OPERATIONS RESEARCH
INFORMS 2016

 | ISBN 978-0-9843378-9-7

Multiagent Systems Modeling

Sanmay Das
Washington University in St. Louis, sanmay@wustl.edu

Abstract A multiagent system is one where multiple autonomous agents with potentially different goals interact. Viewing agents through the computational lens provides a powerful, yet principled method for understanding the behaviors of complex systems, including economic and financial markets, online social networks, etc. In this tutorial, I discuss general principles for such modeling, best practices for handling the simplicity/complexity tradeoff, and present examples of predictive and useful models.

Keywords Agent-based modeling; financial markets; prediction markets; bounded rationality; zero-intelligence traders; market-making

1. Introduction

The term complex systems is used in many ways. For our purposes, we view complex systems as ones whose interesting and important properties arise as a result of interactions between multiple agents who may have different objectives and different capabilities. The emergent behavior of a complex system is determined by several factors, including the behaviors of the agents, the rules of engagement, the network of agent interactions, and the structures of information available to the agents. As such, there is an enormous range of choices that can be made by the human modeling the system.

There are several different disciplines that engage in modeling complex systems, and they have evolved different conventions for how to make these choices. In this tutorial, I will discuss the emerging norms for modeling complex systems from the perspective of agent design, using the "computational lens." In many ways, this can be viewed as falling between the strict norms of agent rationality in economics and the statistical mechanics approach from physics. In doing so, I will focus on two major application areas: modeling of economic and financial markets, and of social networks.

This tutorial is structured as follows. I start by describing typical characteristics of the methodologies for modeling adopted by mainstream economics and the physics-based approach. I then turn to the agent-based or computational approach, and describe the ways in which modelers design the agents that participate in these systems, including the design of boundedly rational agents and agents designed using psychological theories or heuristics. After that, I give an example of this approach and how it has been used in modeling financial and prediction markets. I conclude with some thoughts on best practices for agent-based modeling, and, in particular, on handling the simplicity/complexity tradeoff.

2. Modeling Methodologies From Economics and Physics

Despite the increasing prevalence of "data-first" approaches to science and engineering, there are many domains in which it is still critical to start from a generative model of a system. Perhaps the most important reason for this is that generative models can allow us to determine causal mechanisms, thus enabling us to ask counterfactual or policy-related questions. Another issue is that models simplify the real-world, and thus allow us, as humans, to interpret it more efficiently. Two of the major methodologies for modeling complex systems come from the disciplines of economics and physics, so it is useful to briefly characterize and compare these methodologies.

2.1. Mainstream economics

Mainstream economics models are typically characterized by partial or general equilibrium analysis. The key idea is that each agent is simultaneously solving an optimization problem, and the main features of the system emerge from the simultaneous solutions of these problems, mediated by the market mechanisms through which agents interact. The tools of optimization and game theory are applied to setting up and solving these models.

Economists place a high value on analytical tractability and model parsimony. They tend to simplify models until agent behavior and interactions can be reduced to a set of equations that provide intuition to the person analyzing the system. In the words of Brian Arthur, "conventional economic theory chooses not to study the unfolding of the patterns its agents create, but rather to simplify its questions in order to seek analytical solutions" [3]. Models are simplified until the decision problems faced by agents are *easy*. This is not meant to be pejorative. Some of these easy decision problems require sophisticated mathematics to solve, and the elegance of solutions can be a factor in determining how valuable the work is considered, but, at the end of the day, few if any mainstream economics models involve an agent having to solve a truly complex decision or learning problem that would need to be solved *algorithmically* rather than analytically.

While this convention enables high quality standards in terms of the elegance and intuition afforded by the models, it can also be unfortunate, because in the real economic systems we are trying to model, agents often do solve exactly such problems. Even if we restrict ourselves entirely to the worlds of finance and banking, where behavior is already more circumscribed (and therefore, in some meaningful way "easier" than decision-making that relates to living the rest of life), examples abound. Banks and trading firms hire hundreds of specialists in machine learning and physics in order to develop sophisticated trading algorithms. Analysts and traders who figured out optimal structures for converting collateral pools into securitized tranches were handsomely rewarded in the run-up to the financial crisis of 2007-08. Decisions on granting credit and managing credit lines are made using complex machine-learning models that predict delinquency.

Not only is it that agents are solving complex problems, these tend to be powerful agents, with substantial influence on social and economic outcomes. In such situations, it is only natural that the rules of the game are designed, and regularly re-designed, with these complex agents in mind. At some point, the reductive focus on analytical expressibility and solutions becomes counter-productive, and it is important to understand how complexity affects outcomes.

2.2. Physics

In a well-known review of work in complex network modeling, Albert and Barabasi write "[Physics], a major beneficiary of reductionism, has developed an arsenal of successful tools for predicting the behavior of a system as a whole from the properties of its constituents" [2]. It is undoubtedly true that physics and physicists have had a significant amount to say about various complex systems, ranging from communication, social, and ecological networks to financial markets. There are several reasons for this, both cultural and related to the standard toolbox of the discipline. Physics brings with it both the focus on reducing complex behavior to simple interactions, as well as a focus on statistical and empirical regularities. The central insight in this regard is that aggregate behaviors of systems of interacting agents are often sufficiently determined by the network structure of interactions and the specific "rules of encounter" that they can be modeled as the outcomes of random, or near-random systems of interacting particles.

This level of granularity in analysis has led to some interesting insights, including models that are used to understand fundamentals of network structure, like the Watts-Strogatz model [51] which can explain the "small world" phenomenon observed in various real-world networks, and the preferential attachment model of Barabasi and Albert [4] which helps to explain the scale-free nature of networks like the world-wide web or citations. The Ising model, a model of ferromagnetism in statistical mechanics, has been used to study opinion formation in social networks [1]. In modeling financial markets, physicists have been responsible for the development and elucidation of market properties based on "zero intelligence" models of trader behavior, originally from the economics literature (more on these below) [27, 22]. Physicists have also had success in documenting and turning attention to various empirical facts, especially the so-called "stylized facts" of stock markets, like fat-tailed return distributions and clustering of volatility [23].

While they have been useful and have pushed progress in several interesting directions, it is worth keeping in mind that these physics-based models of complex systems necessarily do not provide a coherent theoretical foundation for how to think about them when the actors are humans or institutions with significant power in shaping outcomes. That is the point at which modeling the behavior of the agents in the system as random becomes too much of a simplification. In financial and economic markets, there are many cases where we need to explicitly reason about the incentives and behavior of market participants in order to understand observed outcomes. In the design of networks like the Internet, there can be many engineering decisions that go into choices that govern the topologies of the resulting networks, and this can yield significantly different properties in, say fault tolerance, than one would expect from preferential attachment networks or other random processes [9].

3. The Agent-Based Approach to Modeling

Agent-based models are characterized by heterogeneity in one or all of agent information, preferences, and strategies. Agent-based models could go all the way from the "simple" models of mainstream economics to representing in detail the full complexity of the real world (after all, each human being is one agent, and together we make up the global economy and society).

Thomas Schelling developed a classic model of segregation that can be thought of as one of the first generation of agent-based models in the 1970s, showing how segregation can arise easily from a slight preference for a minimum number of similar neighbors [44]. Epstein and Axtell's work on Sugarscape in the 1990s builds agent based models of general artificial societies, where agents may have to compete for resources (hence the "sugar") but can also interact with each other and the world in more complex ways [20]. Depending on the specific application, Sugarscape can be used to model different social systems of interest. In finance, agent-based models have been used to model different types of trader behavior in order to try and explain several empirical puzzles that cannot be resolved by more traditional models [34].

Over the last three decades, the increase in computational power and programming proficiency have both contributed to a proliferation of agent-based models in different domains. Agent-based modeling has provided a big tent. It has room for all kinds of agent descriptions, ranging from very simple agents to agents informed by psychology and behavioral economics to complex, deliberative agents that are attempt to optimize sophisticated performance metrics. The flexibility that these models afford is a great attraction, but at the same time leads to a question that cuts straight to the heart of the enterprise: What is the right agent description in these models?

3.1. Types of agents

While it is difficult to characterize all the different ways in which agents can be modeled, there are several types of agent models that have appeared often in the literature. We discuss them under three main categories: (1) simple agents that act probabilistically, (2) agents that are, or strive to be, rational, and (3) agents that are also self-interested, but designed from a different perspective than decision-theoretic rationality, including agents that mimic human behaviors, and agents that use heuristics.

3.1.1. Stochastic agents As mentioned above, one of the insights that has helped physicists to build useful models of markets is that aggregate behaviors of systems of interacting agents are often sufficiently determined by the network structure of interactions and the specific "rules of encounter" that they can be modeled as the outcomes of random, or near-random systems of interacting particles. These ideas are not unique to physics. In fact, the seminal paper of Gode and Sunder on how markets with zero-intelligence trading agents can come close to allocative efficiency was from mainstream economics [27]. In their model of zero-intelligence (ZI) trading, machine traders subject to a budget constraint choose i.i.d. bid or offer prices for single shares, with the prices being drawn uniformly at random across the entire bounded, feasible range of trading prices. The key insight is that the double auction mechanism, in combination with the agents' budgets, actually constrains market outcomes regardless of agent behavior.

This idea has spurred a significant amount of follow-on work across computer science, econophysics, and economics. Cliff and Bruten showed that the ZI model can lead to very different price properties than would be expected in equilibrium when market demand and supply are asymmetric [13]. They proposed a family of "ZI plus (ZIP)" traders, that use a simple form of machine learning to adapt their desired profit margin over time and demonstrated that continuous double auction markets with these ZIP traders demonstrate transaction price properties more consistent with data from human traders. The ZIP model itself has added to the toolbox of agent designers, especially in trading agent competitions [40].

Farmer et al use a zero-intelligence model of trading to explain properties of order data from the London Stock Exchange like the bid-ask spread [22]. In their model, impatient traders take liquidity by placing market orders. while patient traders place limit orders uniformly on the semi-infinite interval defined by the best price on the other side of the market.[1] Again, it is interesting that such a simple model of price choice can explain a large portion of the dynamics of bid and ask prices.

Othman adapts zero-intelligence agents to the prediction market setting [42], and Das uses them in microstructure models that are focused on market making [16, 17]. In both of these cases the generative model is somewhat different. Agents come to the market with a belief w about the true value of the instrument being traded. Facing market bid and ask prices b and a, traders could either choose to buy immediately (a market order) at a (if $w > a$), sell immediately at b (if $w < b$), or place a limit buy order at some price $b' < w$ (if $b' \geq a$ this effectively becomes a market buy), or a limit sell at $a' > w$ (if $a' \leq b$ this effectively becomes a market sell). Depending on the support of prices (for prediction markets, this would be $[0, 1]$), the distributions of limit prices could come from different families, typical choices are uniform or Gaussian. This type of generative model also requires specification of the population distribution of beliefs (such a distribution can also be backed out for models that do not start from agent beliefs). Othman [42] and Brahma et al [7] use families of Beta distributions in the prediction market setting, while Das uses point distributions [16] as well as Gaussians [17].

[1] Technically, they work in the log space, so an arriving trader who wishes to place a buy order would choose one in the interval $[-\infty, a)$ where a is the log of the current best ask. They use a normalization trick to get around the problem of choosing a log-price uniformly at random over an infinite interval.

It is also worth pointing out a connection to a second thread of literature. In classic models of market microstructure, there are often "noise traders" present – these are traders who are modeled as trading for non-information reasons, and therefore typically randomly buy or sell. There are two classic justifications for why these are reasonable for modeling the world. The first is to model long-term investors: those who buy into the stock market, for example, in order to save for retirement, and then sell in order to finance consumption in retirement. The second is simply as a model of those who may try to beat the market without sophisticated information processing capabilities – day traders or technical traders, for example (although technical traders may be doing something systematically different from noise trading, as I discuss below). In truth, the introduction of noise traders can also be thought of as a technical device that allows models of market microstructure to get around the no-trade theorem of Milgrom and Stokey, which says that in an equilibrium where the information structures are common knowledge, adverse selection would result in no trade between differentially informed rational agents [39].

Stochastic agents also often do not need to be made explicit. Consider the literature on network cascades. The main question is how decisions made by nodes in a network affect the overall dynamics of the network [50]. Models of information [6], product adoption [19], influence [33], or disease spread [21] in networks, for example, typically consider stochastic transmission where a node is influenced/infected/etc. with some probability as a function of the number of its neighbors that are. Some very influential (no pun intended) models in this area have been the independent cascade and threshold models. In independent cascade models, once a node u becomes activated (infected, adopts a product, etc.), it can be considered contagious, and then it has one chance of influencing each of its neighbors v, each with some probability p. Therefore, in a round t, if k of v's neighbors have *just* become active, the probability that v will be activated in this round is $1-(1-p)^k$. One important factor is that each node only gets one chance to influence each of its neighbors. In threshold models, by contrast, each node u has a threshold θ_u, and becomes active when a fraction θ_u of its (weighted) neighborhood is active. Active nodes remain active in perpetuity, but the network will converge to a steady state.[2]

Note that in all of these network models, the behavior of individual agents (nodes) is very simple, stochastic behavior, but it can still lead to interesting results at the network level. In a similar fashion, the the idea that certain parts of large, complex economic systems like financial markets can be abstracted away and treated *as if* random or close-to-random in some well-defined way is a powerful one.

3.1.2. (Boundedly) rational agents The agents of mainstream economics theory, *homo economicus* as they are sometimes called, are optimizers. They come endowed with utility functions, and maximize their utility functions subject to constraints on budgets, information, and so on, as well as expectations about strategic interactions.

The "full" rationality found in models from mainstream economics is both implausible and intractable for all but the simplest decision problems. Herb Simon identified the problem more than a half century ago [45].

> Broadly stated, the task is to replace the global rationality of economic man with a kind of rational behavior that is compatible with the access to information and the computational capacities that are actually possessed by organisms, including man, in the kinds of environments in which such organisms exist. One is tempted to turn to the literature of psychology for the answer. Psychologists have certainly been concerned with rational behavior, particularly in their interest in learning phenomena. But the distance is so great between our present psychological knowledge of the learning and choice processes and the kinds of knowledge needed for economic and administrative theory that a marking stone placed halfway between might help travelers from both directions to keep to their courses.

[2] These are classic models that have influenced the whole line of literature. See Kempe et al [33] for some more historical context.

Foundations for bounded rationality How should we go about modeling boundedly rational agents? An alternative comes from the agent-based approach to AI, which involves designing agents that "do the right thing" [43]. The "right thing" in this context means that agents should take actions that maximize their expected utility (or probability of achieving their goals). Ideally, an agent should be perfectly rational. Unfortunately, there is really no such thing as a perfectly rational agent in the world. As Russell says, "physical mechanisms take time to process information and select actions, hence the behavior of real agents cannot immediately reflect changes in the environment and will generally be suboptimal." Calculative rationality, the ability to compute the perfectly rational action in principle, given sufficient time and computational resources, is not a useful notion, because agents that act in the world have physical constraints on when they need to choose their actions. We are left to contemplate other options for agent design.

Russell proposes two other options for a goal for agent design – *metalevel rationality* and *bounded optimality.* Metalevel rationality involves reasoning about the costs of reasoning. An agent that is rational at the metalevel "selects computations according to their expected utility" [43]. This is what Conlisk refers to as deliberation cost [14], and both Russell and Conlisk make the explicit connection to the analogous value of information. Conlisk argues strongly for incorporating deliberation cost into optimization problems that arise in economics, saying that human computation is a scarce resource, and economics is by definition the study of the allocation of scarce resources. He suggests that instead of optimally solving an optimization problem P, a decision-maker should solve an augmented problem $F(P)$ in which the cost of deliberation is taken into account. The problem, as he realizes, is that it is also costly to reason about $F(P)$, and, therefore, theoretically at least, one should reason about $F(F(P)), F(F(F(P))), \ldots$ This infinite regress is almost always ignored in the literature that does take deliberation cost into account, assuming that $F(P)$ is, in some sense, a "good enough" approximation. Economics is not alone in this myopic consideration of deliberation cost. In the AI literature, the tradition of studying deliberation cost (or the essentially equivalent concept which Russell calls the value of computation) as part of solving a decision problem dates back to at least the work of Eric Horvitz [29], and almost all the algorithms that have been developed have used myopically optimal metareasoning at the first level, or shown bounds in very particular instances. The history of metalevel rationality in the AI literature is more that of a useful tool for solving certain kinds of problems (especially in the development of anytime algorithms) than as a formal specification for intelligent agents. This is mostly because of the infinite regress problem described above as Russell writes (of the first metalevel, or the problem $F(P)$ in Conlisk's notation), "perfect rationality at the metalevel is unattainable and calculative rationality at the metalevel is useless."

This leaves us with Russell's last, and most appealing, candidate — bounded optimality, first defined by Horvitz as "the optimization of [utility] given a set of assumptions about expected problems and constraints on resources" [29]. Russell says that bounded optimality involves stepping "outside the agent" and specifying that the *agent program* be rational, rather than every single agent decision. An agent's decision procedure is boundedly optimal if the expected utility of an action selected by the procedure is at least as high as that of the action selected by any decision procedure subject to the same resource bounds in the same environment. This puts the onus of rationality on the agent designer, but that seems to make more sense as a one-time optimization problem than requiring the same of the agent for every optimization problem it faces.

While it may be hard to achieve, with the burden of rationality shifted from the agent to the agent designer, bounded optimality seems to be the right goal. If one agent algorithm can be shown to perform better than another given the same beliefs about the set of problems the agent may face and the same computational resources, the first algorithm should be the one used. Of course, this can be problematic because it is not clear where it must perform better. What is the set of problems the agent can be expected to solve? What if the prior beliefs of the agents are completely wrong?

A possible solution is to analyze these questions from the perspective of the agent designer. It is good scientific practice to use the best available methods when attempting to solve a problem, so we should hold boundedly optimal algorithms to the standard of reality. The performance of algorithms should be tested on problems that are as realistic as possible. If the algorithm performs well, it is a good agent description to use. Many publications in AI and machine learning are focused on designing the best performing algorithms for various problems, ranging from object recognition to trading competitions. Detailed specifications and even implementations of these algorithms are usually publicly available, making it easier to use these agent descriptions in practice.

Examples of bounded rationality in modeling While they do not explicitly make this connection, there are a few examples of how the notion of bounded optimality, in the sense of designing the best agent program you can, is implicitly incorporated into the design of agents used in modeling. One is the empirical game theoretic analysis (EGTA) program of research of Michael Wellman and his group [52]. EGTA is applied in scenarios where agents are playing traditional non-cooperative games which are too large to be solved using conventional methods. The EGTA process involves using a restricted set of strategies for each agent and then searching for an equilibrium in that space of restricted strategies using various techniques from simulation and game reduction [32]. Given the space of strategies available to agents, EGTA can find approximate equilibria. The process helps gain traction in understanding the outcomes of complex systems of agents when the game is known, but very difficult to solve. It has been used in interesting ways in analyzing, among other things, competition between different types of financial exchanges [48], the welfare effects of market making [49], and the formation of credit networks [15]. It is interesting to note that the restriction to a specified space of strategies is usually important to have across all agents to have any hope of solving the system, but may also be a sensible restriction for an individual agent given limited reasoning capabilities and the increased difficulty of performing strategic reasoning across game-theoretic scenarios.

Another line of literature relates to trading agent competitions, where the goal is to design trading agents to compete in specific tasks that the tournament designers set forth [46, 53], but then the resulting trading agents can be used as libraries of intelligent agents that were designed to perform well at the specified task. Trading agent competitions in the AI community have spurred the design of sophisticated agents for all kinds of interesting market games, including energy trading, ad auctions, and supply chain management. There have been several studies that have actually used the agents that performed well in these games in building models to understand market properties better [31].

3.1.3. Behavioral and heuristic agents The heuristics and biases program made famous by Kahneman and Tversky studies actual human behavior and how it deviates from the norms of rational choice [47]. This program is not in any way prescriptive, as it mainly focuses on cataloging deviations from the presumed normative laws of classical decision theory. Thus, this program does not provide any suitable definitions for intelligence. While some models of prospect-theoretic agents have been successful at explaining real-world phenomena, for example the equity premium [5], prospect theory has not really emerged as a reasonable guideline for agent-based modeling, perhaps because it explicitly requires deviations from known rational behavior in modeling, and the main reason to do so would be to study precisely that deviation and its effects.

Another approach from the literature of cognitive science is the use of "satisficing" heuristics in the tradition of Simon [45], who introduced the notion that human decision-makers do not exhaustively search over the space of outcomes to choose the best decision, but instead stop as soon as they see an outcome that is above some satisfactory threshold "aspiration level." Conlisk cites various papers in the economics literature that start from Simon's notion of bounded rationality, and claims that, within economics, "the spirit of the idea is pervasive" [14]. In cognitive science and psychology, Gigerenzer and others have recently popularized the use of "fast and frugal" heuristics and algorithms as the natural successor to satisficing. Gigerenzer and Goldstein state their view of heuristics as being "ecologically rational" (capable of exploiting structures of information present in the environment) while nevertheless violating classical norms of rationality [24]. They have a program to design computational models of such heuristics, which are "fast, frugal and simple enough to operate effectively when time, knowledge, and computational might are limited" while making it quite clear that they do not agree with the view of heuristics as "imperfect versions of optimal statistical procedures too complicated for ordinary minds to carry out." They reject the Kahneman-Tversky program because it maintains the normative nature of classical decision theory.

It is interesting that Gigerenzer and Goldstein held a simulated contest between a satisficing algorithm and "rational" inference procedures, and found that the satisficing procedure matched or outperformed more sophisticated statistical algorithms. The fact that a simple algorithm performs very well on a possibly complex task is not surprising in itself, but what is very clear is that if it can be expected to perform better than a sophisticated statistical algorithm, it must either be a better inference procedure or have superior prior information encoded within it in the context of the environment in which it is tested. As agent designers, if we had to design an agent that solves a given range of problems, and had access to the information that a satisficing heuristic was the best known algorithm for that range of problems, it would be silly not to use that algorithm in the agent, all else being equal. The problem with fast and frugal heuristics as a program for agent design is the loose definition of what constitutes a satisfactory outcome, or of what kinds of decision-making methods are "ecologically rational" in the language of Gigerenzer and Goldstein. How do we know that one heuristic is actually better than another?

Some heuristic agents arise naturally within a discipline and it makes sense to incorporate them in agent-based models. For example, consider "technical traders" in finance. "Technical analysis" is in itself a broad umbrella term that covers a wide range of ways of trading, but the essential idea is to use rules based on time series patterns in order to predict future behavior of a stock. These rules typically tend to have some kind of intuitive geometrical interpretation, and therefore technical analysts are also sometimes called "chartists." They are essentially heuristics for making decisions on when to buy and sell financial instruments. There has been widespread disparagement of technical trading in the academic finance literature [38], partly because the efficient markets hypothesis should preclude the ability to make profits based on publicly available information. Nevertheless, there has also been research showing that there are inefficiencies in market prices, and some of these may be exploitable using technical trading rules [37], and indeed, some of these technical trading rules are not that different from econometric methods used to identify market inefficiencies. The adaptive markets hypothesis offers a possible explanation, saying that markets are not necessarily efficient, but instead prices reflect information that is determined by the environment of trading for that asset and the types of agents engaging in trade [36]. This justification is also related to the notion of ecological rationality mentioned above. In any case, the prevalence of these techniques in real markets, and the evidence that they may have some success can be compelling reasons to include them in agent-based models.

3.2. Critiques and directions

While agent-based models have had many successes, it is hard to conclusively demonstrate that the ability of these models to replicate real-world phenomena comes from some real insight into the system being modeled, as opposed to from the huge increase in the degrees of freedom available to the designer of the agent-based model. Along with the problem of allowing models to overfit and easily replicate past phenomena without leading to better *predictive* power, the availability of many parameters makes such models less transparent; even without the intention to do so, many assumptions can end up being implicit, rather than exposed, as they are when agent decision-making can be written down in a few equations. As such, agent-based models have received criticism as being too "ad hoc" because of the dual problem of providing less intuition and model flexibility allowing for easier fits to past data and events.

At the same time, there is a clear need for new models of economic and social systems that capture complexity in ways that are crucial for understanding, regulating, and designing these systems. For example, the Institute For New Economic Thinking at Oxford calls for "a more realistic view...[that] sees the economy as a dynamic, complex, evolving, network of interacting, heterogeneous individuals and institutions who don't always behave rationally and have limited information, but nonetheless learn, are innovative, and evolve over time."[3] How can we arrive at such a view? Any such paradigm for modeling must:

- Start from the basic principles that agents respond to incentives, and attempt to make good decisions, subject to constraints on information and reasoning capabilities. Thus, the classic economics notions of equilibrium, in which agents are responding as best as they can to prices, institutional rules, and each other, are the right lens through which to try and understand system-wide outcomes.
- Recognize the importance of modeling the complexity of real-world systems, rather than abstracting important elements away in the search for analytical solutions, unique equilibria, and so on. This requires engagement with dynamics, multiple and approximate equilibria, and approximate solutions to decision problems.
- When appropriate, use insights from the "zero intelligence" and econophysics literatures to abstract certain actors in the economy into pseudo-representative agents that can have significant influence on markets through their aggregate power, but may be constrained by institutional design or regulation into acting in ways that make them easy to represent stochastically, instead of strategically.
- Use a principled approach towards formulating the decision problems that agents must solve and what algorithms they will use to solve them. This approach must be immune to the common criticism of agent-based modeling (and bounded rationality research) summarized by John Conlisk as follows — "Without the discipline of optimizing models, economic theory would degenerate into a hodge podge of ad hoc hypotheses which cover every fact but which lack overall cohesion and scientific refutability" [14].

I now turn to discussing a line of work that is representative of this proposed approach to modeling.

4. An Example Application: Modeling Financial and Prediction Markets

In this section, I illustrate the approach to modeling through an example related to financial and prediction markets. The central question in market microstructure research is how the rules that govern trading affect price formation, and hence the information aggregation and dissemination properties of markets. Many markets, for example the NYSE and NASDAQ, have historically used market-makers. Market makers solve the chicken-and-egg problem of liquidity – by always being willing to take the other side of a trade, they encourage other agents to participate in the marketplace, thus creating even more liquidity [41].

[3] http://www.inet.ox.ac.uk/why-we-need-new-economics (March 14, 2016)

4.1. The Glosten-Milgrom model

One of the gold-standard models of market-making is that of Glosten and Milgrom, who characterize how market makers should set bid and ask prices (the prices they are willing to buy and sell shares at, respectively) under conditions of asymmetric information, where the agents they are trading with may be better informed than them [26]. They characterize how adverse selection affects the spread between the bid and ask prices, and therefore how information can affect price movements, returns, etc. The model is a classic mainstream economics model, with a constrained market-maker, informed traders with superior information about asset value, and noise traders. The market-maker must set prices at any time without knowing if she is trading with an informed or uninformed trader, and the development of the model is such that the market-maker could attempt to optimize different objectives, although Glosten and Milgrom focus on zero-profit pricing, the expected competitive outcome.

Assume that there is a single instrument, with true value V, traded in the market. Trading opportunities arrive sequentially in discrete time, and at time t, the market maker sets bid and ask prices b_t and a_t respectively for one unit of the instrument. At each time t, a single trader with signal $w_t = V + \epsilon_t$ arrives, and buys one unit from the market maker if $w_t > a_t$, sells one unit to the market maker if $w_t < b_t$, and does not trade otherwise. One of the core insights of Glosten and Milgrom is that a zero-expected-profit market-maker should set the ask price to the expected value of the security conditional on receiving a buy order, and the bid price to the expected value of the security conditional on receiving a sell order (note that this restricts the information available to traders so that, in order for these prices to exist, the market maker must always be able to trade off expected losses from more informed traders with expected gains from less informed traders, getting around the no-trade theorem). Thus, the zero-profit equilibrium prices (without consideration of transaction or inventory costs) are $b_t = E[V|w_t < b_t]$ and $a_t = E[V|w_t > a_t]$. In general, this can be a complex fixed-point equation, but, as we shall see below, there are cases where the prices can be efficiently computed.

4.2. Market making in more complex models

Glosten and Milgrom's model is an elegant and powerful framework for thinking about the asymmetric information component of the spread, but they only solve for extremely simple models of trader behavior and information. What if we were to enrich the space of trader models? The first thing is that clean analytical solutions to the bid and ask prices are no longer easy to find. Nevertheless, by constructing an algorithmic implementation of a market-maker (which must be experimentally validated to achieve reasonable performance in terms of its goals) we can derive insights into the properties of more realistic markets. Das [16, 17] examines models with jumps in the underlying true value (so now the true value process is stochastic and v_t is indexed by time). In these models, the market-maker is again trying to set zero-profit prices in expectation given her uncertainty about the true value, and there are two different models for traders. In one of the models, there are perfectly informed and perfectly uninformed traders. A perfectly informed trader at time t receives signal $w_t = v_t$, while a perfectly uninformed trader can be modeled as receiving a signal $w_t = v_t + \epsilon_t$ where $\epsilon_t = \pm C$ with probability 0.5 that $\epsilon_t = +C$ and probability 0.5 that $\epsilon_t = -C$. While C can be calibrated to spreads to yield different probabilities of trading, in these models C is set to a large enough magnitude that uninformed traders always trade in one direction or the other. In the other model, the traders all receive signal realizations from the same (Gaussian) distribution, and $w_t = v_t + \epsilon_t$ where $\epsilon_t \sim \mathcal{N}(0, \sigma_0^2)$. This model can replicate some interesting features of real markets, for example, it produces a two-regime behavior with significant heterogeneity of information and large spreads following a jump in the true value (for example, shortly before an earnings announcement, when information may have leaked to insiders), which gets resolved by the market-maker's information updates to allow the market to return to a regime of relative homogeneity of information and small spreads [16].

The complexity of agents can be extended both on the market-maker side and the trader side, and even more questions can be explored with richer models. For example, an important practical and regulatory consideration in markets that employ market-makers has historically been whether to employ a single monopolist specialist (as the NYSE did) or allow multiple market-makers to compete (as the NASDAQ did). Market-makers would like to maximize their own profit (which is why they get into the game), but this may conflict with the system-wide goal of increasing liquidity. The conventional argument would be that competition between market-makers would therefore lead to better social outcomes, since it would drive the spread down. This was not observed empirically in the case of the NYSE vs. NASDAQ over the years [30, 12]. In fact, the NYSE used to promote the benefits of a monopolist specialist for "maintaining a fair and orderly market" in the face of market shocks [25].

4.2.1. An optimal monopolist An extension of the model described above is to develop a more sophisticated market-making model in order to study the dynamic profit-maximization problem of a monopolist market-maker. The problem here is that the dynamic optimization problem of the market-maker is over the state space of her *beliefs* $p_t(v)$ about the true value, which is difficult to solve without using approximations except in the simplest cases. Das and Magdon-Ismail's solution is to design an approximate dynamic programming algorithm based on moment-matching approximations for the market-maker's posterior beliefs [18]. They operate under the assumption that the true value v is fixed, and all traders are noisy informed traders receiving signals $w_t = v + \epsilon_t$ where $\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$, and σ_ϵ is common knowledge. The optimal strategy is theoretically given by the Bellman equation for the value functional V

$$V(p_t;\pi) = \mathbb{E}[r_0|p_t, b_t^\pi(p_t), a_t^\pi(p_t)] + \gamma\mathbb{E}[V(p_{t+1};\pi)|p_t, b_t^\pi(p_t), a_t^\pi(p_t)]$$

where the state p_t, is a *function* encapsulating the market maker's belief, (b_t, a_t) (the bid-ask pair) is the action, and π is a policy mapping states to actions.

In order to solve this problem in an efficient manner, Das and Magdon-Ismail develop an algorithm that maintains a Gaussian belief over the true value v, with parameters μ_t for the mean and σ_t for the variance collapsing the true posterior back to a Gaussian by matching the first and second moments. They assume symmetric bid and ask prices around the mean belief, and this allows the decision problem to be framed in terms of a single variable (the variance of the market-maker's belief, σ^2). They prove that σ^2 is monotonically decreasing, which enables an efficient, single-sweep approximate solution for the dynamic programming problem defined above.

The algorithm leads to new phenomenological insights: Das and Magdon-Ismail find that the monopolist profit maximizer provides *more* liquidity than perfectly competitive market-makers in periods of extreme uncertainty about the true value of an asset. The mechanism is that the monopolist is willing to absorb some losses in those periods by maintaining a tighter spread in order to more quickly learn the true valuation, information that can be exploited for more profits later (see Figure 1). Competitive market-makers would never be willing to absorb a loss in any period, since the informational benefits would accrue to everyone, and they would not be guaranteed the ability to profit from the information themselves. This was the first model to demonstrate the intuition that a learning benefit may be the reason why monopolists can provide better markets in some conditions. Of course there is a tradeoff, because the expected spreads of the monopolist-mediated market are higher than those in the competitive case in times of relative tranquility (when there is not much uncertainty or asymmetry of information).

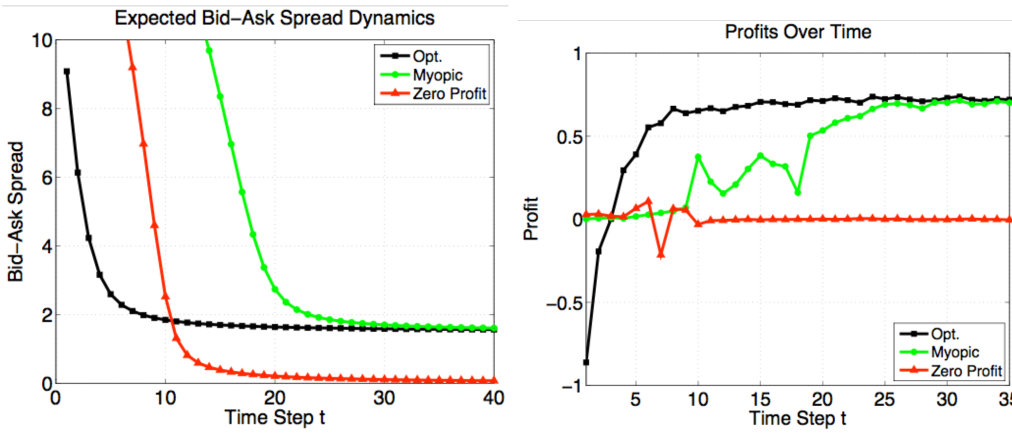

FIGURE 1. The average bid-ask spread over time (left) and average profits (right) in simulations of a dynamically optimal monopolist, a competitive market-maker, and a myopically optimal market-maker [from [18]]. The optimal dynamic monopolist is willing to absorb losses early on by maintaining an artificially low spread in order to more quickly learn the true value and profit from that information.

4.2.2. More complex traders The other natural extension mentioned above is to more complex trader models. Brahma et al set out to design a market-maker based on the models discussed above and apply it to real-world *prediction markets* [7]. For our purposes here, a prediction market can simply be thought of as a market on an asset that pays off an amount in $[0,1]$ upon the realization of some event. For example, an election prediction market could pay off 1 if a Democrat is elected president, and a 0 otherwise (or it could pay off the proportion of the national vote won by the Democrat). The dominant method for market-making in prediction markets is to use *scoring rule*-based market makers [28, 11], which set prices using a potential function (sometimes called a cost function) on the inventory position. While these algorithms have many nice properties, they are necessarily loss-making whenever the prediction market is actually doing a useful job of prediction [7]. One advantage of methods based on the finance literature is that they may be able to not make losses in expectation, which would be very important in real-money markets.

Brahma et al make the above work of Das and Magdon-Ismail practical by including modules in the algorithm for dealing with trade size consideration and for automatic detection of possible jumps in the true value based on likelihoods of observed sequences of trades [7]. Trade size is dealt with through the heuristic technique of treating orders as sums of “mini batches” of orders. In a pure dealer-mediated market, the dealer has to quote a price for any quantity. When traders only demand single units, the bid and ask prices can be computed as described above, but say an order of size Q is received at time t, when the market maker’s belief is parameterized by μ_t and σ_t. There is a “mini batch size” parameter α in the market-making algorithm. Then this order is treated as $\lceil Q/\alpha \rceil$ independent orders with sizes $\alpha_1, \alpha_2, \ldots, \alpha_k$ (all α, except possibly the last). The α parameter is a heuristic designed to deal with the fact that the orders are not actually independent, but the size of the order still may convey information. The market maker has to quote a single price. She internally simulates the arrival of k orders, performing k state updates and computations to arrive at k different prices for the mini orders, and then quotes the volume weighted average price (VWAP) of these k orders. However, note that if the quoted price is accepted by the trader, the *actual* moment-matching belief update is based only on the VWAP price, not the sequence of simulated prices.

Automatic detection of possible jumps is a more difficult issue. The intuition is that. following a jump, we would expect to see a more imbalanced series of trades. Brahma et al define a consistency index $C(\text{history})$, intended to measure the relative likelihood of the recent history of trades under the current uncertainty level of the market maker, as opposed to under a higher uncertainty. The market maker keeps track of a fixed window of previous trades. Each trade is associated with values z^+ and z^-, which are the upper and lower bounds of the w, the signal the trader must have received in order to be willing to execute that trade given the prices. The probability of a sequence of trades over a window of size W, is then:

$$L(\mu,\sigma)=\int_{-\infty}^{\infty}N\big(v,\mu,\sigma\big).\prod_{i=1}^{W}\Big(\Phi\big(z_i^+,v,\sigma_\epsilon\big)-\Phi\big(z_i^-,v,\sigma_\epsilon\big)\Big)\,dv$$

The consistency index enables a relative comparison with the same probability computed at twice the uncertainty.

$$C(\text{history})=L(\mu_t,2\sigma_t)-L(\mu_t,\sigma_t)$$

If $C>0$, the market maker sets $\sigma_{t+1}=2\sigma_t$, automatically also widening future spreads.

Since they set out to design a practical market-making algorithm, Brahma et al also test and evaluate it with many different collections of trading agents, as well as with human subject experiments. For the trading agent experiments, the true price process is modeled as follows. The price p_t evolves in discrete time $0,\ldots,T$. p_0 is sampled uniformly at random on $[0,1]$. At time t, the price jumps with some probability ρ. If the price does not jump, $p_t=p_{t-1}$, otherwise, p_t is sampled from a normal distribution with mean p_{t-1} and fixed variance σ^2_{jump}. If $p_t\geq 1$ or $p_t\leq 0$ after a jump, the instrument is assumed to liquidate immediately at 1 or 0 respectively. There are two types of trading agents, fundamentals traders and technical traders, which exemplify the principles of bounded optimality and heuristic agents.

Fundamentals traders The fundamentals traders are expected linear-utility maximizers, subject to a constraint on the variance of the final utility. This involves having to maintain a belief on the current true value and estimate the distribution of the final utility based on that belief. Let us consider the second problem first.

Perfectly accurate computation is again difficult, so suppose instead that the trader uses a Gaussian approximation for future beliefs. For any given point belief $\hat{p}_t$, given a utility function $u(S,C,\hat{p}_T)$ (defined over the liquidation value), a number of shares S of the security, and cash C, the expected final utility can then be computed as:

$$\begin{aligned}U_{\text{point}}(t,\hat{p}_t,S,C)=\sum_{k=0}^{T-t}\Bigg[&B(k;T-t,1/T)\\&\left(\prod_{j=1}^{k-1}\Phi(0<x<1;\hat{\mu}_j,\hat{\sigma}_j)\int_0^1 d\hat{p}_T\phi(x=\hat{p}_T;\hat{\mu}_k,\hat{\sigma}_k)u(S,C,\hat{p}_T)\right.\\&\left.+H(\hat{p}_t,k)u(S,C,1)+L(\hat{p}_t,k)u(S,C,0)\right)\Bigg]\end{aligned}$$

Here $B(k;n,p)$ is the binomial probability mass function with the jump probability, $\phi(x;\mu,\sigma)$ is the normal PDF, and $\Phi(x;\mu,\sigma)$ is the normal CDF. The functions H and L are the probability that $p_T=1.0$ (high) and $p_T=0.0$ (low) respectively given that there are k jumps. Any of the jumps might cause the stock to liquidate, in which case the remaining jumps do not actually happen.

$$H(\hat{p}_t,k)=\sum_{t=1}^{k}\left[\Phi(x\geq 1;\hat{\mu}_t,\hat{\sigma}_t)\prod_{j=1}^{t-1}\Phi(0<x<1;\hat{\mu}_j,\hat{\sigma}_j)\right]$$

$$L(\hat{p}_t, k) = \sum_{t=1}^{k} \left[\Phi(x \le 1; \hat{\mu}_t, \hat{\sigma}_t) \prod_{j=1}^{t-1} \Phi(0 < x < 1; \hat{\mu}_j, \hat{\sigma}_j) \right]$$

The remaining question is how to estimate the future beliefs, parameterized by $\hat{\mu}_t, \hat{\sigma}_t$ at time t. Each jump adds a normally distributed random variable, then truncates, leading to a non-Gaussian distribution after two jumps. However, the Gaussian approximation is quite good for small numbers of jumps, and large numbers of jumps become increasingly improbable. Given this approximation, $\hat{\mu}_j$ and $\hat{\sigma}_j$ for the estimated mean and standard deviation after j jumps can be computed as the mean and standard deviation of a truncated Gaussian plus a normal with mean 0 and standard deviation σ_{jump}:

$$\hat{\mu}_j = \hat{\mu}_{j-1} + \hat{\sigma}_{j-1} \frac{\phi(\frac{-\hat{\mu}_{j-1}}{\hat{\sigma}_{j-1}}) - \phi(\frac{1-\hat{\mu}_{j-1}}{\hat{\sigma}_{j-1}})}{\Phi(\frac{1-\hat{\mu}_{j-1}}{\hat{\sigma}_{j-1}}) - \Phi(\frac{-\hat{\mu}_{j-1}}{\hat{\sigma}_{j-1}})}$$

$$\hat{\sigma}_j^2 = \sigma_{\text{jump}}^2 + \hat{\sigma}_{j-1}^2 \left[1 + \frac{\frac{-\hat{\mu}_{j-1}}{\hat{\sigma}_{j-1}} \phi(\frac{-\hat{\mu}_{j-1}}{\hat{\sigma}_{j-1}}) - \frac{1-\hat{\mu}_{j-1}}{\hat{\sigma}_{j-1}} \phi(\frac{1-\hat{\mu}_{j-1}}{\hat{\sigma}_{j-1}})}{\Phi(\frac{1-\hat{\mu}_{j-1}}{\hat{\sigma}_{j-1}}) - \Phi(\frac{-\hat{\mu}_{j-1}}{\hat{\sigma}_{j-1}})} - \left(\frac{\phi(\frac{-\hat{\mu}_{j-1}}{\hat{\sigma}_{j-1}}) - \phi(\frac{1-\hat{\mu}_{j-1}}{\hat{\sigma}_{j-1}})}{\Phi(\frac{1-\hat{\mu}_{j-1}}{\hat{\sigma}_{j-1}}) - \Phi(\frac{-\hat{\mu}_{j-1}}{\hat{\sigma}_{j-1}})} \right)^2 \right]$$

Now, we turn to describing the information process and current beliefs of the fundamentals traders. These traders are weakly informed in the following manner: they are aware of the true value stochastic process according to which p evolves, and at each period t, they each receive an independent sample from a Bernoulli trial with success probability p_t. The model includes two types of fundamentals traders: both attempt to infer the true value from the history of samples they have received (and prior knowledge of the jump process), but one trades only on the basis of its belief about the true value as inferred from its private history of samples, while the other one also incorporates knowledge of the market price under a weak rational expectations model.

The "private information" trader (or "beta trader") maintains a Beta distribution as her belief over possible values of p_t. She assumes that there has not been a jump in the past W time periods. Even if this assumption is violated, the distribution shifts naturally as old information leaves the window and is replaced by post-jump samples. With k successes in the past $n \le W$ trials, expected final utility is given by

$$U_{\text{beta}}(t, n, k, S, C) = \int_0^1 d\hat{p}_t \text{Beta}(\hat{p}_t; k+1, n-k+1) U_{\text{point}}(t, \hat{p}_t, S, C) \tag{1}$$

$\text{Beta}(x; \alpha, \beta)$ is the beta PDF. The beta trader attempts to maximize U_{beta} by picking the number of shares S which maximizes $U_{\text{beta}}(t, n, k, S, C)$, taking transaction fees for exchanging cash and shares into account, subject to the variance constraint, which can be taken into account by computing $\mathbb{E}_{\hat{p}_T}[(u(S, C, \hat{p}_T))^2]$.

The private information trader is a pure fundamentals trader. By incorporating price history, we can get a more accurate estimate of the true underlying market value than from private information alone. The rational expectations trader is very similar to the market maker described above, the main difference being that a trader does not need to quote prices, and does not receive information about cancels. Instead, the ask price in the market maker's inference algorithm is set to an order's execution price. Belief updates and adaptivity to shocks then follow the same process. The mean μ and variance σ^2 of the belief is then combined with the beta distribution from private information as follows:

$$U_{\text{RE}}(t, n, k, S, C) = \frac{1}{Z} \int_0^1 d\hat{p}_t \text{Beta}(\hat{p}_t; k+1, n-k+1) \phi(\hat{p}_t; \mu, \sigma) U_{\text{point}}(t, \hat{p}_t, S, C)$$

$Z = \int_0^1 dx \text{Beta}(x; k+1, n-k+1)\phi(x; \mu, \sigma)$ normalizes the combined probability distribution. σ is lower bounded for these "rational expectations" traders, to avoid drowning out private information. One interpretation of this combined distribution is that the agent samples from each separately until drawing the same value from both. This favors values which are likely according to both the agent's private information and the belief inferred from market prices.

Technical traders The technical traders in the model are based on prior literature [8]. These are commonly known as the "moving average" rule and "trading range break." The moving average rule compares a short-term moving average (say over one day) and a long-term moving average (say over 200 days) of a stock price, and buys when the short-term moving average rises above the long-term one, and sells when it falls below the long-term one. In their model, Brahma et al use two-state moving average traders that keep track of a long (30 trade) and short (10 trade) execution price average, L_t and S_t respectively at time t. If the trader is in the "low" state and $S_t > (1+\alpha)L_t$, it buys and moves to the "high" state; if the trader is in the "high" state and $S_t < (1-\alpha)L_t$, it sells and moves to the "low" state. α determines a margin, which helps avoid responding to noise.

The trading range break-out rule is to buy when the price rises above a "resistance" level (a local maximum) and sell when the price falls below a "support" level (a local minimum). In a fixed window (20 trades), traders find the highest execution price $P_{\max}$ and the lowest $P_{\min}$. If the most recent execution price is more than $(1+\alpha)P_{\max}$, a trader with a trading opportunity buys. If it is less than $(1-\alpha)P_{\min}$, the trader sells. α again discourages responding to noise,

4.2.3. Experimental validation Including technical traders like the above in the model, creates more realistic conditions for the boundedly rational market-making algorithm to be validated against. Brahma et al are able to achieve many of their desiderata for a market-maker (low spreads, low losses, good liquidity provision and market quality) in these environments with fundamentals and technical traders, as well as in a series of short human subject trading experiments.[4]

In another, related piece of work, Chakraborty et al use this algorithm to make markets in a long-running "instructor rating" use of prediction markets in which all the participating traders are humans (competing for prizes based on their virtual currency winnings) [10]. Again, the market-making algorithm is shown to be successful at achieving its objectives of not making loss, providing good liquidity and price stability, etc. The parallel evolution of agent design and modeling here is particularly relevant. The development of this market-making algorithm came about from both the motivation to better model markets and to use these algorithms to design better markets. The paradigm was to design the best market-making algorithm possible for the specific problem being studied, but as the space of trader behavior became richer, this necessitated more approximations and algorithmic innovation. Thus, the algorithm eventually had a credible claim to being the best market-making algorithm known for prediction markets, and it was based on the same ideas that were used to model outcomes in financial markets.

4.3. Policy implications

It is worth returning briefly to the notion of agent-based models as generative models that can be useful for studying policy questions and counterfactuals. A detailed but principled model of the kind above is useful for pursuing such questions. If we were to design or redesign a stock exchange or a new market of some other kind today, these models can be directly used. Given expectations about the relative amounts of time the market would

[4] Another interesting example of a game that shares the same ideology of design is the TAC Market Design Competition (CAT) [40]. In this competition, participants aim to design better mechanisms to maximize a score (a combination of profit, market share and transaction success rate) when traders are drawn from a known population of different types, including ZI, ZIP, and other traders.

spend in high vs. low information asymmetry states, and social or exchange-directed goals for what the spread should look like, a model like ours allows for a better characterization of the relative benefits and costs of monopolists. Tools like this can predict policy effects and allow us to examine possible futures in a principled way – for example, if prediction markets were to be deregulated and real-money prediction markets entered the US financial ecosystem, how could we ensure sufficient liquidity, and what rules should govern trading around times of high uncertainty? The market-making algorithms have also recently been used to study questions about the efficiency and equity of continuous double auctions as opposed to frequent call markets [35], the latter of which have been proposed as a policy response to some of the deleterious consequences of high-frequency trading.

5. Concluding Thoughts: Best Practices

Getting a multiagent model "right" is more a matter of good scientific and engineering practice rather than strict adherence to some pre-defined set of rules for what constitutes an "allowable" model, whether those rules have to do with agent rationality or otherwise. The key tradeoff is really about the right level of complexity for the modeler to choose in designing both the model and the agents that participate. The purpose of a model is not to create an exact simulation of reality, but instead to offer generative or causal explanations of observed behavior, allow one to pose counterfactual questions, evaluate policy choices, and predict future behavior. As models become more complex, they become more capable of approximating reality. At the same time, the addition of parameters and design decisions can make them more fragile, they may overfit the past, and they may provide less intuition and interpretability. There is no single solution to this – while mainstream economics may still be a step too far on the simple side of the equation, this does not mean we should (even if we could) build models in which all agents (for example) must actually solve the problems solved by high-frequency trading platforms on specialized, low-latency hardware. The central question of machine learning today is that of how much to regularize models in order to achieve the highest out-of-sample predictive performance, and this may offer lessons in thinking about the "right" level of complexity for models. At the end of the day, we want our models to lie somewhere on an efficient frontier of interpretability, predictivity, and robustness. We can achieve this with careful attention to detail in modeling and an appreciation of the evolution of the area.

Acknowledgments

I am grateful to Mithun Chakraborty and Aparna Gupta for comments on a draft of this article, and to the NSF for research support through awards 1414452 and 1527037.

References

[1] Mohammad Hadi Afrasiabi, Roch Guerin, and Santosh S Venkatesh. Opinion formation in ising networks. In *Information Theory and Applications Workshop (ITA), 2013*, pages 1–10. IEEE, 2013.

[2] Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47, 2002.

[3] W. Brian Arthur. Complexity and the economy. *Science*, 284(5411):107–109, 1999.

[4] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

[5] Nicholas Barberis, Ming Huang, and Tano Santos. Prospect theory and asset prices. *The Quarterly Journal of Economics*, 116(1):1–53, 2001.

[6] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, pages 992–1026, 1992.

[7] Aseem Brahma, Mithun Chakraborty, Sanmay Das, Allen Lavoie, and Malik Magdon-Ismail. A Bayesian market maker. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 215–232, 2012.

[8] William Brock, Josef Lakonishok, and Blake LeBaron. Simple technical trading rules and the stochastic properties of stock returns. *The Journal of Finance*, 47(5):pp. 1731–1764, 1992.

[9] Jean M Carlson and John Doyle. Highly optimized tolerance: Robustness and design in complex systems. *Physical Review Letters*, 84(11):2529, 2000.

[10] Mithun Chakraborty, Sanmay Das, Allen Lavoie, Malik Magdon-Ismail, and Yonatan Naamad. Instructor rating markets. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 159–165, 2013.

[11] Y. Chen and D.M. Pennock. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*, pages 49–56. Citeseer, 2007.

[12] William G. Christie and Paul H. Schulz. Why do NASDAQ market makers avoid odd-eighth quotes? *Journal of Finance*, 49(5):1813–1840, 1994.

[13] Dave Cliff and Janet Bruten. Zero not enough: On the lower limit of agent intelligence for continuous double auction markets. *HP Laboratories Technical Report HPL*, 1997.

[14] John Conlisk. Why bounded rationality? *Journal of Economic Literature*, 34(2):669–700, 1996.

[15] Pranav Dandekar, Ashish Goel, Michael P Wellman, and Bryce Wiedenbeck. Strategic formation of credit networks. *ACM Transactions on Internet Technology (TOIT)*, 15(1):3, 2015.

[16] Sanmay Das. A learning market-maker in the Glosten-Milgrom model. *Quantitative Finance*, 5(2):169–180, April 2005.

[17] Sanmay Das. The effects of market-making on price dynamics. In *Proceedings of the International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 887–894, May 2008.

[18] Sanmay Das and Malik Magdon-Ismail. Adapting to a market shock: Optimal sequential market-making. In *Advances in Neural Information Processing Systems (NIPS)*, pages 361–368, 2008.

[19] Pedro Domingos and Matt Richardson. Mining the network value of customers. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 57–66. ACM, 2001.

[20] Joshua M Epstein and Robert Axtell. *Growing Artificial Societies: Social Science From the Bottom Up*. Brookings Institution Press, 1996.

[21] Stephen Eubank, Hasan Guclu, VS Anil Kumar, Madhav V Marathe, Aravind Srinivasan, Zoltan Toroczkai, and Nan Wang. Modelling disease outbreaks in realistic urban social networks. *Nature*, 429(6988):180–184, 2004.

[22] J. Doyne Farmer, Paolo Patelli, and Ilija I. Zovko. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences*, 102(6):2254–2259, 2005.

[23] J. Doyne Farmer, Eric Smith, and Martin Shubik. Is economics the next physical science? *Physics Today*, 58(9):37–42, 2005.

[24] Gerd Gigerenzer and Daniel G. Goldstein. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4):650–669, 1996.

[25] Lawrence R. Glosten. Insider trading, liquidity, and the role of the monopolist specialist. *Journal of Business*, 62(2):211–235, 1989.

[26] Lawrence R. Glosten and Paul R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14:71–100, 1985.

[27] Dhananjay K Gode and Shyam Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.

[28] R. Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.

[29] Eric Horvitz. Reasoning about beliefs and actions under computational resource constraints. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 429–444, 1987.

[30] Roger D. Huang and Hans R. Stoll. Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE. *Journal of Financial Economics*, 41(3):313–357, 1996.

[31] Patrick R Jordan, Christopher Kiekintveld, and Michael P Wellman. Empirical game-theoretic analysis of the TAC supply chain game. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 193–200. ACM, 2007.

[32] Patrick R Jordan, Yevgeniy Vorobeychik, and Michael P Wellman. Searching for approximate equilibria in empirical games. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1063–1070. International Foundation for Autonomous Agents and Multiagent Systems, 2008.

[33] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146. ACM, 2003.

[34] Blake LeBaron. Agent-based computational finance. *Handbook of Computational Economics*, 2:1187–1233, 2006.

[35] Zhuoshu Li and Sanmay Das. An agent-based model of competition between financial exchanges: Can frequent call mechanisms drive trade away from CDAs? In *Proceedings of the International Joint Conference on Autonomous Agents and Multi-Agent Systems*, 2016. To appear.

[36] Andrew W Lo. Reconciling efficient markets with behavioral finance: The adaptive markets hypothesis. *Journal of Investment Consulting*, 7(2):21–44, 2005.

[37] Andrew W Lo, Harry Mamaysky, and Jiang Wang. Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation. *The Journal of Finance*, 55(4):1705–1770, 2000.

[38] Burton Gordon Malkiel. *A Random Walk Down Wall Street*. WW Norton & Company, 1999.

[39] P. Milgrom and N. Stokey. Information, Trade and Common Knowledge. *Journal of Economic Theory*, 26(1):17–27, 1982.

[40] Jinzhong Niu, Kai Cai, Simon Parsons, Peter McBurney, and Enrico H. Gerding. What the 2007 TAC market design game tells us about effective auction mechanisms. *Autonomous Agents and Multi-Agent Systems*, 21(2):172–203, 2010.

[41] Maureen O'Hara. *Market Microstructure Theory*. Blackwell, Malden, MA, 1995.

[42] Abraham Othman. Zero-intelligence agents in prediction markets. In *Proceedings of the International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 879–886, 2008.

[43] Stuart J. Russell. Rationality and intelligence. *Artificial Intelligence*, 94(1):57–77, 1997.

[44] Thomas C Schelling. Dynamic models of segregation. *Journal of Mathematical Sociology*, 1(2):143–186, 1971.

[45] Herbert A. Simon. A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69:99–118, February 1955.

[46] Peter Stone and Amy Greenwald. The first international trading agent competition: Autonomous bidding agents. *Electronic Commerce Research*, 5(2):229–265, April 2005.

[47] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.

[48] Elaine Wah, Dylan R. Hurd, and Michael P Wellman. Strategic market choice: Frequent call markets vs. continuous double auctions for fast and slow traders. In *Proceedings of the Conference on Auctions, Market Mechanisms and Their Applications (AMMA)*, 2015. To appear.

[49] Elaine Wah and Michael P Wellman. Welfare effects of market making in continuous double auctions. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 57–66. International Foundation for Autonomous Agents and Multiagent Systems, 2015.

[50] Duncan J Watts. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99(9):5766–5771, 2002.

[51] Duncan J Watts and Steven H Strogatz. Collective dynamics of small-worldnetworks. *Nature*, 393(6684):440–442, 1998.

[52] Michael P Wellman. Methods for empirical game-theoretic analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 21, pages 1552–1555, 2006.

[53] Michael P. Wellman, Amy Greenwald, and Peter Stone. *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition (Intelligent Robotics and Autonomous Agents)*. The MIT Press, August 2007.