
Toward Agent-Based Models for Investment

J. Doyne Farmer
McKinsey Professor
Santa Fe Institute
Santa Fe, New Mexico

Although agent-based models are not yet ready for practical investment application, they can yield powerful insights about market behavior, particularly in regard to the “second order” inefficiencies that create profit-making opportunities. When practical use of agent-based models becomes possible (perhaps within the next five years), their effectiveness will cause securities prices to change.

As far as I know, no one currently uses agent-based models for investment. In the context of this presentation, an agent-based model involves a model for price formation and a model for agent behavior, including the models’ information inputs and how the inputs behave, which could involve learning. From an investor’s point of view, agent-based models represent a new direction that may or may not result in more practical investment models. From an academic’s point of view, agent-based models are fascinating new tools that make it possible to investigate problems previously out of reach. The models presented here do not provide practical tools, but they do offer a taste of several new directions that may or may not turn out to have practical value in the future.

One of the things that attracted me to financial prediction is the possibility of forecasting the behavior of people. People have free will—you never know what they will do. Of course, we make predictions about people in our day-to-day lives. We all know people whose behavior is pretty consistent and whose responses are predictable. But making quantitative predictions about people in general is a real challenge, and little success has been apparent so far.

Because prices involve people, predicting prices is different from predicting the weather. When the weatherman says it is likely to rain, this prediction has no effect on whether it will rain. In contrast, a

statement by George Soros that gold prices are going to rise can have a significant effect on the price of gold. Predicting prices is more like telling people’s fortunes. If a fortune-teller says you will take a long trip and have a terrible accident and you have any faith in her forecast, you will stay home. If you do stay home and no terrible accident occurs, was she right? The point is that forecasts regarding humans, who have free will and can alter their behavior, can be self-invalidating. This point is the idea underlying the theory of efficient markets: If someone discovers a pattern, it diminishes as it is exploited. If enough people discover the pattern, it will disappear.

Informational efficiency (the lack of arbitrage) is a central idea of modern finance. Anybody who has spent time building market models knows that markets are arbitrage efficient at some level. Building models that make good market predictions is not easy. But markets cannot be perfectly efficient; otherwise, why would so many smart people waste their time investing in them? Milton Friedman originally pointed out the paradox of efficient markets: If rational speculation makes markets efficient, then because the market is efficient, no profits can be made and all the rational speculators should leave, thus causing the market to revert to an inefficient state. The theory of efficient markets is inherently self-contradictory. Finding a self-consistent theory to replace it is an important problem that demands a solution. As a physicist, I would state the situation as follows: At first order, markets are arbitrage efficient, but at second order, they are not. For markets to function, there

Author’s note: This presentation is partially based on research done in collaboration with John Geanakoplos, Shareen Joshi, and Paul Melby.

have to be small inefficiencies to motivate investors. These small inefficiencies are what sustain the ecology of the market. All the action, all the interesting behavior of markets, is in the second order. This point is clearly understood by practitioners and should also be understood by economic theorists.

Finance is like any other industry. Profit-making opportunities, or inefficiencies, are created by the evolving needs of customers. Specialization is driven by economies of scale in digesting and understanding information. One must regard the market as an ecology of highly specialized, heterogeneous agents, linked together by their relationships to each other and the flows of money that these relationships imply.

The story I am going to tell is about how the demands of investors create inefficiencies. In this case, inefficiencies arise because investors demand liquidity, thus driving a need for market makers. Market makers have risk aversion, which means that once they acquire positions, they have to off-load them. But risk aversion means that their actions are somewhat predictable. For example, if they off-load their positions gradually, they cause trends in prices. Trends in prices are exploited by technical traders, who analyze patterns in past prices and trade accordingly. Thus, we see how the demand for liquidity can sustain a population of technical traders, even though the technical traders are not providing liquidity directly. Money flows from liquidity demanders to technical traders. This pattern does not mean the market is not pretty efficient; it just means that this approximate efficiency comes about only through a web of interactions among heterogeneous players, who are all part of an interconnected market ecology. But I am getting ahead of my story.

Fundamental vs. Empirical Models

My own experience at Prediction Company and elsewhere illustrates the difference between agent-based models, which are a fundamental approach to modeling, and empirical models, which are a strictly practical, data-driven approach.¹ Norman Packard, Jim McGill, and I were the senior founding partners of Prediction Company. When we founded Prediction Company, we did not know anything about financial markets. We were experts in time-series modeling, which is an empirical approach to building models for data that are ordered in time. I came to be involved in time-series forecasting because of my earlier expe-

rience in the field of chaotic dynamics. I developed nonlinear time-series models that could take advantage of low-dimensional chaos and, in some cases, give improved forecasts for such problems as fluid turbulence, ice ages, and sunspot cycles.

Prediction Company uses a purely empirical approach to market forecasting. The models are not based on a fundamental, reasoned understanding of how markets work; instead, they are based on a search for patterns in historical data. The fundamental premise is that patterns that existed in the past will repeat themselves in the future. These patterns are usually rather subtle. The market is pretty efficient in that extremely few persistent simple patterns exist. It is not quite so efficient with respect to complex patterns, and if you look carefully and comprehensively, you may find things that very few have seen before. The big problem is that once you allow for complex patterns, the number of possibilities is enormous and careful statistical testing becomes essential.

Feature extraction is a key aspect of Prediction Company's models. In the extraordinarily noisy environment of the financial world, statistical modeling methods that blindly generate models directly from data do not work well. The problems are too much noise, the limited size of datasets, and the fact that techniques have a hard time finding patterns that are real and persistent rather than ephemeral random fluctuations. Throwing away distracting information is crucial.

The same principles apply in many other problems that require cognition, such as image recognition. Suppose you want to design a machine that can recognize objects and interpret visual images. You might naively try to do it by operating directly on the pixels of an image. But this approach has not been successful. We know the brain does not work this way, as was dramatically illustrated in a famous set of experiments by two neuroscientists, Hubel and Wiesel, who studied vision in spider monkeys. They gave the monkeys a variety of different visual stimuli and recorded the output of selected neurons in the visual cortex of the monkeys. They showed that the image processing in the cortex happens in stages. The first several stages are for pre-processing, to encode the information from the retina into features that are easier to interpret. Neurons are specialized to respond to features, such as edges, lines, motion in a specific direction, or differences between the right and left eye. It appears that these features evolved to facilitate the process of decomposing an image into understandable features and properties, such as plants, animals, trees, and objects and their movement with respect to each other. Irrelevant information is discarded—indeed, most of the bits of information in the image are thrown away so that the

¹Prediction Company acts as a fund advisor and technology development group for a proprietary, fully automated trading operation of Warburg Dillon-Read, which is an affiliate of the Union Bank of Switzerland.

brain can focus all its attention on the small subset of information that is likely to be useful.

Finding the right signal-processing feature detectors to make sense out of a financial time series is not easy, because we have no a priori knowledge of what these features are. Evolution at Prediction Company has proceeded on a timescale of years rather than millennia, driven by researchers formulating and testing new feature detectors and accumulating lore that is brought to each successive modeling problem. This area is where we spend most of our time and where we believe we gain most of our edge. We originally based our search for good features on the technical-trading literature, which has evolved over roughly a hundred years. As you might expect, we found that most technical-trading rules are useless by themselves, but we do not use them by themselves. Instead, we just regard them as pre-processing steps, signal filters that separate the data into features, such as "15-day trend" or "maximum under-5-day smoothing." Although standard technical-trading rules formed the starting point for our original models, the later models also involve many signal filters (not based on technical-trading rules) that we discovered through a lot of hard work and trial and error.

This approach involves a great deal of human intervention, and it is not cut-and-dried science but rather an art form with statistical validation as its critic. But it is still almost entirely empirical. Although we may structure our models based on ideas about what makes markets tick, the models are still mainly based on fitting these signal-processing filters to data in concert with standard function-approximation algorithms, such as linear regression and artificial neural nets. Ideas about how markets work drive the choice of the inputs and shape the way the data are presented to the learning algorithms. But we cannot say that the models reflect a fundamental understanding of what people are doing and how the market works. Rather, the trading models are merely complex, unconscious, stimulus-response machines. (The trading is completely automated, with no human intervention.)

One can make an analogy to the prediction of a physical phenomenon, such as the trajectory of a baseball. Suppose an engineer is given the task of training a robot to catch baseballs. There are two approaches: empirical or fundamental. The empirical approach mimics how the ball is caught. A human learns to catch baseballs through a process of trial and error, by seeing a lot of baseballs get hit and forming a largely unconscious set of mental maps, with empirical rules for forecasting the trajectory of the ball and responding to it. The engineer can mimic this

process by installing appropriate learning algorithms and training the robot. Alternatively, the engineer can use her fundamental knowledge of physics. Without ever watching a ball get hit, she can construct a model for the ball's motion in terms of the laws of gravity, air friction, and kinematics. She could use physics to develop an algorithm for the robot to parse the image from its video camera eyes, estimate the velocity and position of the ball, and use Newton's laws to tell the robot how to forecast where the ball will go and how to move in order to catch it.

Prediction Company's approach is like the human approach, and agent-based modeling is like the approach of the engineer discussed above. This analogy is apt because it illustrates the strengths and weaknesses of each method. In a situation with lots of historical data and little fundamental understanding, the empirical method is superior. But it has the disadvantage that when no historical data are available (e.g., if conditions in the market are fundamentally different from what they have been in the past), the empirical method will fail. The fundamental method is harder to develop, but once developed, it can be powerful in dealing with situations that are outside normal experience. For predicting the trajectories of baseballs, I doubt that the fundamental method can beat the empirical method. But for other tasks, such as predicting the motion of satellites, the fundamental method is clearly superior. (In fact, for the baseball problem, I am willing to bet that most engineers would opt for a mixture of the two methods.)

The pros and cons of these two methods were brought home to me by an experience I had while I was in graduate school. Norman Packard and I developed a system for beating roulette by treating it as a physical system, that is, measuring the position and velocity of the ball after it is released by the croupier and taking advantage of the fact that bets can be placed until a few seconds before the ball exits the track. We built concealable computers (which was quite difficult in 1976) and used them to time the motion of the ball and predict its likely landing point. The early version of the hardware and software implemented expressions based solely on physics. As a result, the program had the advantage of being extremely small and could be fit in the 3,000 bytes (not megabytes or kilobytes!) of memory that were available to store the program and the 128 bytes of RAM that were available for scratch space. The disadvantage was that the parameters of the physical solution were not orthogonal, so it was difficult to distinguish the effect of a change in one parameter (the fall-off velocity of the ball) from a change in the other parameter (the deceleration rate of the ball).

Only I and one other physicist were able to reliably master the manual fitting procedure.

A few years later, when we had a lot more memory and much faster computers, we revisited the problem and solved it with an empirical method. We formatted the data in a natural way that made use of the fact that the wheel is a circle, and we fit a cubic spline to the data. This change made the algorithm easy to use. No knowledge of physics was required. But this approach was not as accurate as the fundamental method when the parameters were properly fitted—because the empirical solution was not as constrained.

Agent-based models have the potential to give us a more fundamental understanding of how markets work and better tools to predict them. They will have the advantage over empirical models in that they can be applied to previously unseen market conditions. Like the development of physics, achieving this understanding will be hard fought and will take a long time to develop. Success will depend on developing good agent models, either by hard-wiring the behaviors of the agents or, more likely, by mimicking their learning process. The advent of practical agent-based models is still at least several years away. Of course, as soon as profitable models are developed, they will begin to make markets more efficient and eventually will invalidate their own forecasts. In the far future, financial speculation likely will be an escalating arms race of agent-based models attempting to out-anticipate the aggregate behavior of each other.

Arbitrage Efficiency Takes Time

The possibility of earning a profit in a market environment is largely attributable to the fact that the progression toward market efficiency is slow and that new inefficiencies are created in an ongoing manner by the needs and foibles of market participants. **Figure 1** shows some “declassified” data from one of Prediction Company’s models. We think about our models as being composed of “signals,” which are highly processed versions of a related cluster of data inputs. This figure illustrates the performance of two different signals as a function of time. The y -axis plots the percent correlation of each of these signals with the following forward return of the market; the data are smoothed to make the picture easier to view. Panel A shows that in the mid-1970s, this signal had almost a 15 percent correlation with the future movements of the market (by market, I mean a universe of liquid U.S. stocks hedged with respect to a long list of risk factors). During the 23-year period that is shown, the average strength of the signal slowly declines. This result both illustrates and rebuts arbi-

trage efficiency. On one hand, this decline in correlation can be seen as a vindication of the idea of arbitrage efficiency, illustrating how slowly market efficiency is achieved. On the other hand, it can be interpreted as illustrating how slowly the achievement of market efficiency occurs. This signal has persisted for at least 23 years and is not dead yet. The statistical significance of the signal is so strong that it is easily detected with only a year of data, and anyone who had found this signal a long time in the past could have made a great deal of money by now. If a “risk premium” is associated with this signal, the risk factors must be extremely difficult to detect and are not reflected in the performance of the model, which has overwhelming statistical significance. The point is that although efficiency eventually makes itself felt, it is very slow in arriving.

The second signal, plotted in Panel B, is even more surprising. The signal begins in the early 1980s and grows through time to the present. This increase seems to be inconsistent with arbitrage efficiency. I do not expect this signal to continue to get stronger, but it illustrates that new inefficiencies can be generated—markets do not always become more efficient with the passage of time.


Simplifying Agent-Based Models

After I left Prediction Company, I came to the Santa Fe Institute, where I have returned to my roots in complex systems theory. This time around, however, I am thinking about markets and having fun trying to model representative agents and their interactions. As a physicist, I naturally think about market participants as being similar to molecules, with trading (which necessarily involves an interaction between agents) being similar to a molecular collision. Of course, the problem with this analogy is that these “molecules” are rather complicated; they are people, not ping-pong balls, and their behavior can be quite complex.

The key to good agent-based modeling is to capture the agent behaviors as realistically as possible with a minimum of assumptions. Unfortunately, in many people’s minds, agent-based modeling has come to be associated with highly complicated models in which each agent has his own customized kitchen sink. Such models are so complicated that they are almost as hard to understand as real markets. In order to enhance understanding, and to ever have a hope of using models for prediction (and fitting their parameters to real data), it is important to keep agent-based models simple.

Walrasian Price Formation. Most academic treatments of price formation are based on the Wal-

Figure 1. Strength of Two Proprietary Predictive Signals, 1975–98



rasian mechanism of price formation, which was inspired by the 19th century Paris Bourse. Price setting is done by an auctioneer. The auctioneer does not trade but rather polls those who want to trade for the quantity they are willing to buy or sell at any given price. Transactions take place only when the auctioneer has found a price at which supply equals demand—that is, when the market clears.

This trading mechanism is still used for the London gold fix, where market participants trade twice a day, and the London Metals Exchange, where they trade four times a day, but most markets no longer operate this way. The Walrasian price-setting mechanism has become unpopular for many reasons. The most important reasons are that price setting is a lengthy process and liquidity takes time to accumulate. As a result, transactions are made only at discrete, widely spaced intervals. But there are other reasons as well: Agents do not like to reveal their demand functions because they are worried that others will use this information to “pick them off.” Indeed, most agents do not even know their full demand functions and do not want to take the trouble to formulate them—they just know how much they want to buy or sell and what price they are willing to pay. Another problem is that a Walrasian call market is a difficult environment for liquidity providers and thus may actually result in more-volatile prices.

Price-Formation Rule. For simplicity’s sake, in the agent-based model described next, we have set some requirements on the price-formation rule. Given observed agent behaviors, prices should converge to

equilibrium under appropriate circumstances. Reasonable demand properties should exist: Buying should drive up the price, and selling should drive it down. The inventory of all the agents should remain bounded for reasonable agent behaviors. No agent should have to sacrifice utility. That is, markets are not a philanthropic business. An agent, at least a speculator, is not expected to stay in the market if he is losing money. The price-formation rule should be as simple as possible and, within reason, should match the statistical properties of the data.

Market-Making Model. The following model is simple. It is designed to explore a few basic qualitative properties of markets and prices. In particular, in terms of the effects on price formation, what is the difference between the classic Walrasian call market (with an auctioneer and market clearing) and continuous markets, with competitive market makers? We argue that market making necessarily induces temporal structure in prices. The particular market-making rule we study causes trending. These side effects make it possible to sustain a population of technical traders. The presence of the technical traders increases arbitrage efficiency, but it also causes prices to make larger deviations from “fundamental values.” This model is much too simple to be a good predictor for real prices or to provide a practical model for practitioners. But it provides an example of how simple agent-based models can be built to understand theoretical problems.

This model is based on an order-driven market. Traders place orders—in this case, market orders—

and the auctioneer is replaced by a market maker. Unlike the auctioneer, the market maker provides liquidity by buying and selling. The agents do not know at what price the orders will be filled, so the transactions automatically take place out of equilibrium. Using an order-driven market solves the problems of a Walrasian market, but other problems pop up in their place. Prices can differ from Walrasian prices, which for various sound theoretical reasons should be good prices. As we will see, this difference can mean that prices automatically contain trends.

For the sake of simplicity, we begin with a single asset and assume a single representative market maker. We allow market orders only and synchronize actions so that trades take place at $t, t + 1$, and so on. We use the following price-formation rule:

$$p(t + 1) - p(t) = \frac{1}{\lambda} \left[\sum_{i=1}^N \omega^i(t) - \beta X(t) \right], \quad (1)$$

where

- ω^i = market order of agent i
- p = log price
- $\omega^i(t + 1) = x^i(t + 1) - x^i(t)$
- λ = liquidity
- X = market-maker position
- β = market-maker risk aversion
- x^i = position of agent i

The change in the logarithm of the price is proportional to the sum of the net order imbalance. The first term is the sum of the orders that are placed by the agents at time t , and the second term is the order placed by the market maker, which is always a fraction β of the market maker's current position. The variable X , the market maker's total position, is the total amount that supply and demand are out of balance in the market. The constant of proportionality, $1/\lambda$, can be thought of as liquidity, and it determines the amount that an order of a given size will move the price (a net order of size λ will change the price by a factor of e). Any particular family of agent behaviors can be studied by specifying a set of functions $x^i(t)$ and iterating the resulting equations.

Behavioral Models. Now, I will introduce a simple behavioral model involving four classes of investors: Market makers, which I have already introduced, fundamental (or value) investors, technical traders (or chartists), and liquidity demanders (or noise traders).

Value investors simply take a position based on the perceived fundamental value of an asset. The more underpriced the asset, the larger the position that value investors take in the asset. For the purposes of this model, I do not care how the investors arrive

at what they think is the correct value—just take it as given externally. For convenience, the asset's perceived fundamental value will change over time according to a logarithmic random walk following the equation $v(t + 1) = v(t) + \epsilon(t)$, where $v(t)$ is the logarithm of the value. Thus, we just assume that there is some positive number that the value investor perceives as being associated with the asset, which changes randomly with time. The value investor's position at any given time is given by the function

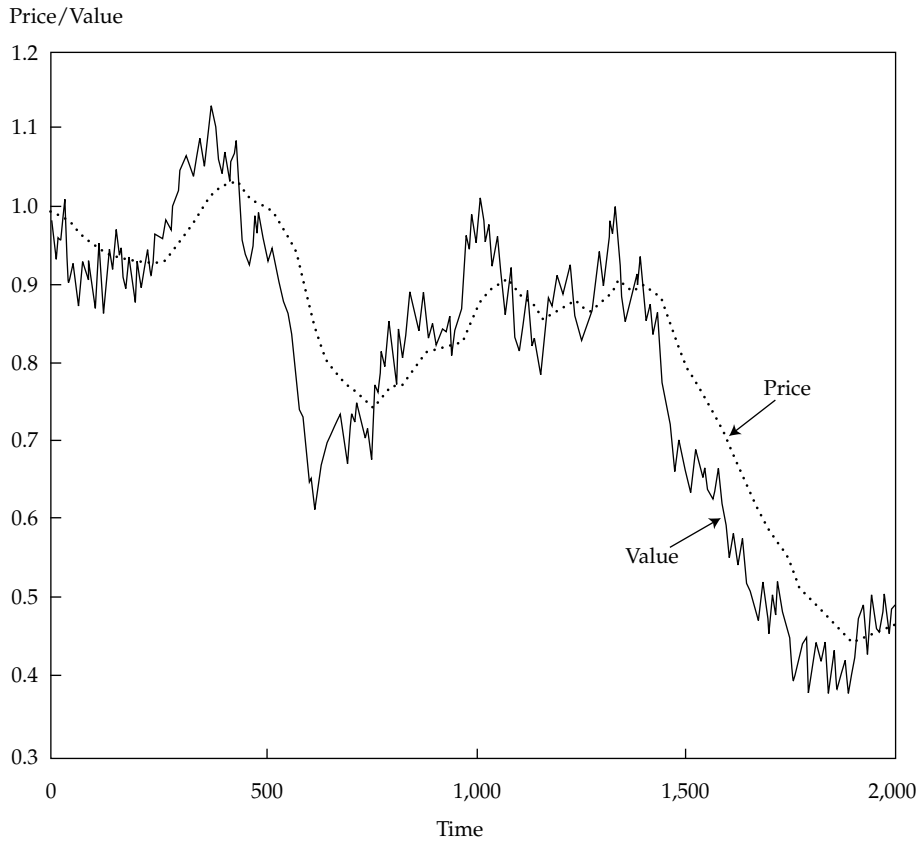
$$x^v(t + 1) = c[v(t) - p(t)].$$

Figure 2 plots price versus time for a scenario in which the only two agents present in the market are the value investor and the market maker. The price roughly tracks the value, oscillating around it but not tracking it perfectly. Figure 2 shows that the model satisfies the criterion that the prices converge to equilibrium under normal conditions. If the value of β is increased, the price will track the value even more closely, and if β is lowered, it will track less closely. If β is lowered all the way to zero, it will not track at all. That is, when $\beta = 0$, the price and the value both just make random walks, but they do so without maintaining any particular relation to each other. We have already learned something: Market-maker risk aversion is important for keeping the price close to its equilibrium value. (If we had used the Walrasian price mechanism, it would always be the case that price = value, so the value is the equilibrium price.) The same relationship is also true of bounded inventories: When $\beta > 0$, inventories are bounded, but when $\beta = 0$, they become unbounded. This pattern is not surprising, because unbounded market-maker inventory means unbounded risk and the β parameter is there to take the market maker's risk aversion into account.

One question that any trader will immediately ask is: Who is making profits? To address this question, we sweep the parameters β and c . (Remember, the parameter c sets the scale of the value investor's trading.) For high values of β and small values of c , the value investor makes a net profit on average, and for the opposite parameters, the market maker makes a net profit on average. The boundary between the two regions is shown in **Figure 3**.

We can probably safely assume that neither player would be happy letting the other make profits at his expense. Thus, if the market maker is making net profits, the value investor will adjust c to increase his profits (and thereby decrease the market maker's profits), and similarly, the market maker can adjust β . This struggle should always result in a solution that is along the boundary curve shown in Figure 3.

Figure 2. Price and Value for a Risk-Neutral Market Maker



But this is not the whole story. If you look closely at Figure 2, you will notice that even though changes in value are completely random, changes in the price are not. This occurs because the market maker’s actions introduce trends into the prices. Once the market maker acquires a position, because of her risk aversion, she has to get rid of it. By selling a fraction β at each time step, she will unload the position a bit at a time. This behavior causes a trend in prices. Any risk-averse behavior on the part of the market maker will result in a temporal structure of some sort in prices.

To provide a better look at the temporal structure in this case, **Figure 4** shows what is called “the power spectrum” of the returns. The power spectrum is the square of the Fourier transform and can be thought of as the amount of “power,” or variance, present at each frequency. If the returns were completely random the power spectrum would be “white,” which would mean all frequencies were equally well represented (i.e., the power spectrum would be constant). Instead, the power spectrum decreases at high frequencies. When $\beta = 0$, the power spectrum is pretty flat throughout most of its range (with a small rise coming from the mean-reverting behavior of the

value investors). When $\beta > 1$, however, it starts out flat and then rolls off. This observation corresponds to the effect of risk-averse market making, which makes prices trend. As β increases, this roll-off

Figure 3. Profitability Constraints

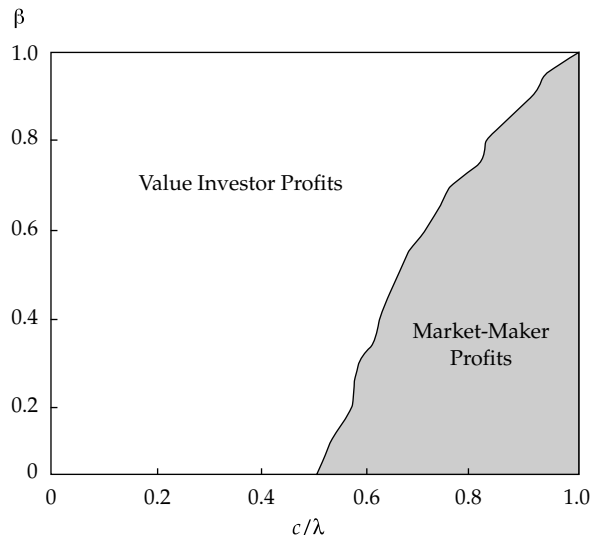



Figure 4. Power Spectrum of Price Returns


becomes sharper, because the power spectrum is nearly white at low frequencies and the timescale of the trending is shorter.

Trend Followers. The fact that the market maker causes a temporal structure in prices creates an opportunity for technical traders. For example, assume this technical trader is a trend follower. A trend follower holds positive positions when the price trend is up and negative positions when the trend is down. For example, consider an exponential moving average:

$$\begin{aligned} x(t) &= cs(t) \\ s(t+1) &= ar(t) + (1-a)s(t), \end{aligned} \quad (2)$$

where

$$\begin{aligned} x(t) &= \text{position} \\ s(t) &= \text{signal} \\ c &= \text{scale} \\ r(t) &= \text{price return} \\ a &= \text{inverse timescale} \end{aligned}$$

This trend-following rule is not necessarily the best possible trading rule to use, but it is sufficient for the trader to extract profits, at least for a reasonable choice of the parameter a , which sets the time scale for trend following, and c , which sets the trading size. In fact, for a reasonable choice of a , by playing with c for the value investor and c for the trend follower, we can find parameter values at which everyone just breaks even. The market is efficient.

By playing with the parameters in the previous example, we showed that a point exists at which everything is efficient. What remains to be shown is whether the market will go there on its own. To find the answer to that question, we must first introduce some new elements into the model—namely, reinvestment and credit limits. Both of these elements have their own side effects.

Credit Limits. Real traders have credit limits (i.e., there is a wealth-dependent limit to the size of the positions they can take). If a trader tries to take a position that is too big with respect to her credit limit and wealth, her position will be capped. If the agent's wealth decreases, she may have to sell some assets to stay within the limit. This limit introduces a nonlinearity into the agent's strategies that has several effects on the price dynamics of the market. It can cause prices to deviate even more from fundamental values and, in extreme cases, can even make prices oscillate on their own.

Another key element is reinvestment. As traders make profits, they change the scale of their trading accordingly. This behavior means that the c parameters are not fixed; rather, they are functions that depend on wealth and credit limits.

When we put all of these elements together, we find that the market does indeed tend toward the point where everything is efficient. This point is somewhat noisy; there are fluctuations around it, but each agent's wealth stays more or less fixed. At this point, there is equilibrium among the market maker, value investor, and trend follower.

Dynamic Market Efficiency. The following example shows how the progression toward efficiency can be much more complicated and can increase rather than decrease volatility. We will now study an entirely different set of agent behaviors. There are two groups. One group consists of seasonal traders that alternately buy and sell. Imagine that they are doing this because they need to (for example, because they are farmers whose crops are harvested at a particular time of year). They are willing to trade to reduce risk or to gain liquidity and do so even if they lose money on average. The other group consists of technical traders. The technical traders are trying to exploit the first group. Unlike the previous technical traders, however, we allow many different technical trading strategies, which we will make up at random, and let the market select the most successful ones.

To keep things simple, assume the technical traders base their strategies on only the signs of the most recent price changes. To construct a particular strategy, list each possible set of signs for past price

changes, which will be called a pattern, and randomly assign either +1 or -1 as the output corresponding to each pattern. The output will be interpreted as being proportional to the position of the strategy. For example, suppose that the pattern is (+,-,+) and the corresponding output is +1. That set of conditions means that if the last three price changes were first up, then down, then up, the strategy would give a "buy" signal. With a buy signal, the agent will invest a fixed fraction of his wealth in the asset (and similarly on the short side for a sell signal).

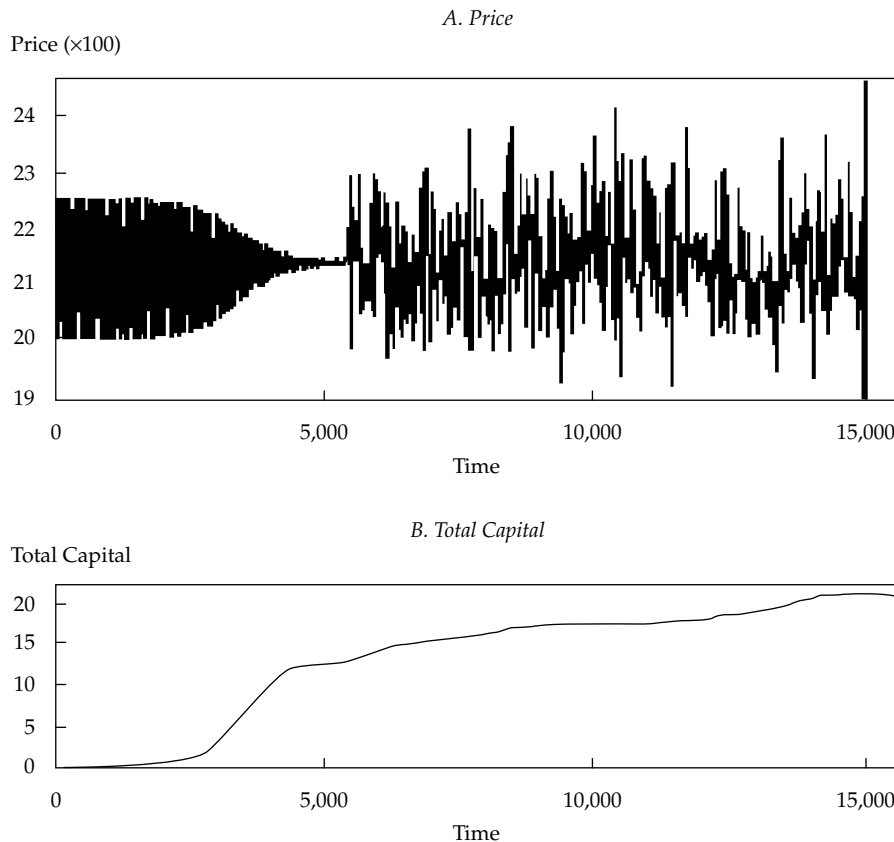
Now, suppose we choose a bunch of traders and randomly assign each of them two strategies (where each strategy is an entire list of responses to every possible pattern). We give all of them a little money to get started and turn them loose.

Figure 5 shows a simulation of this situation, illustrating the long-term evolution toward market efficiency. The beginning of the simulation is on the left side (at time 0). Panel A shows a snapshot of the price every 101 time steps. Panel B shows the total wealth of all the traders. At the beginning, the sea-

sonal traders have almost all the wealth, and the price simply alternates up and down as they trade. After some time passes, the technical traders who have been successful start to make money, and as they make money, they rise in importance. They begin placing bigger orders because their wealth is growing as a result of their trading profits. As this response happens, the oscillating pattern slowly dampens because the price impact of their trades cancels the impact of the seasonal traders.

When time equals 5,000 (the bottleneck seen in Figure 5), a point is reached at which the market seems to be highly efficient. The original oscillating pattern of prices is now virtually gone, cancelled out by the action of the technical traders. Just as they are about to award Nobel Prizes for the theory of market efficiency, however, something odd happens: Prices suddenly go crazy and start oscillating again. This effect is caused by the technical traders, who have gained enough wealth to exert a greater influence on prices than the seasonal traders. Whereas the sequence of signs had previously been always alternating

Figure 5. Long-Term Evolution Toward Market Efficiency



between plus and minus, there are now sometimes two plus signs or two minus signs in a row. This pattern invokes a set of previously unused patterns in the strategies of each agent. Prices start changing in a chaotic, volatile manner again. No longer is the pattern a simple periodic oscillation; rather, the pattern is much more random and is driven by a complex set of interactions among all the agents. The cash flows shift again. Although the market might be getting more efficient, it is certainly not getting less volatile. Strategies that had been unprofitable when the seasonal players were dominating may now become profitable and begin to accumulate wealth.

This scenario happens in the absence of noisy inputs. The simulation is completely deterministic. The statistical properties of prices continue to change, even tens of thousands of iterations later, as the feeding relationships of who is exploiting whom shift around. There is a rich and slowly evolving ecology of agents, with shifting interactions. Market efficiency takes a long time to happen.

Future Research

As I said in the introduction, at this stage, none of these models is of practical use yet. Although we can replicate some of the properties of real prices, we do

not replicate all of them accurately. These models are simple, but we have illustrated certain basic points. We see how the market mechanism matters: Using an order-based market with market makers as liquidity providers can result in patterns in prices that sustain trend followers. When we introduce a large number of different strategies and let the market select them, the path toward efficiency can be surprising because of the dynamics of the strategies interacting with each other.

Agent-based modeling of markets is still in its infancy. I predict that as the models get better, they will begin to be useful for real problems. The work I have presented here is still some distance from that point. If I needed an investment strategy now, I would certainly use the empirical approach. But to really understand how markets work, we have to use agent-based models. Within five years, people may be trading money with agent-based models. Time will tell. The degree to which fortune-tellers affect the future depends on their credibility. Thus, as agent-based models achieve some successes, market participants will start to pay attention, and prices will change in response. Agent-based models are a brand new tool, and the market for them is still far from efficient.