

Simple Forecasts and Paradigm Shifts

Harrison Hong
Princeton University

Jeremy C. Stein
Harvard University and NBER

First draft: March 2003
This draft: September 2003

Abstract: We study the implications of learning in an environment where the true model of the world is a multivariate one, but where agents update only over the class of simple univariate models. If a particular simple model does a poor job of forecasting over a period of time, it is eventually discarded in favor of an alternative—yet equally simple—model that would have done better over the same period. This theory makes several distinctive predictions, which, for concreteness, we develop in a stock-market setting. For example, starting with symmetric and homoskedastic fundamentals, the theory yields forecastable variation in the size of the value/glamour differential, in volatility, and in the skewness of returns. Some of these features mirror familiar accounts of stock-price bubbles.

We are grateful to the National Science Foundation for research support. Thanks also to Patrick Bolton, John Campbell, Glenn Ellison, David Laibson, Sven Rady, Andrei Shleifer, Christopher Sims, Lara Tiedens, Jeff Wurgler and seminar participants at Princeton, the University of Zurich, the University of Lausanne, the Stockholm School of Economics, the Norwegian School of Management and the NBER for helpful comments and suggestions.

I. Introduction

In attempting to make even the most basic kinds of forecasts, we can find ourselves inundated with a staggering amount of potentially relevant raw data. To take a specific example, suppose you are interested in forecasting how General Motors stock will perform over the next year. The first place you might turn is to GM's annual report, which is instantly available online. GM's 2002 10-K filing is more than 180 pages long, and is filled with dozens of tables, as well as a myriad of other facts, footnotes and esoterica. And this is just the beginning. With a few more clicks, it is easy to find countless news stories about GM, assorted analyst reports, and so forth.

How is one to proceed in the face of all this information? Both common sense, as well as a large literature in psychology, suggest that people simplify the forecasting problem by focusing their attention on a small subset of the available data. One powerful way to simplify is with the aid of a theoretical model. A parsimonious model will focus the user's attention on those pieces of information which are deemed to be particularly relevant for the forecast at hand, and will have her disregard the rest.

Of course, it need not be normatively inappropriate for people to use simple models, even exceedingly simple ones. There are several reasons why simplifying can be an optimal strategy. First, there are cognitive costs to encoding and processing the added information required by a more complex model. Second, if the parameters of the model need to be estimated, the parsimony inherent in a simple model improves statistical power: for a given amount of data, one can more precisely estimate the coefficient in a univariate regression than the coefficients in a regression with many right-hand-side variables. So simplicity clearly has its normative virtues. However, a central theme in much of the psychology literature is that

people do something other than just simplifying in an optimal way. Loosely speaking, it seems that rather than having the meta-understanding that the real world is in fact complex, and that simplification is only a strategy to deal with this complexity, *people tend to behave as if their simple models provide an accurate depiction of reality.*¹

Theoretical work in behavioral economics and finance has begun to explore some of the consequences of such normatively-inappropriate simplification. For example, in many recent papers about stock-market trading, investors pay attention to their own signals, and disregard the signals of others, even when these other signals can be inferred from prices. The labels for this type of behavior vary across the papers—sometimes it is called “overconfidence” (in the sense of investors overestimating the relative precision of their own signals); sometimes it is called “bounded rationality” (in the sense that it is cognitively difficult to extract others’ signals from prices); and sometimes it is called “limited attention”. But labels aside, the reduced forms often look quite similar.² The common thread is that, in all cases, agents make forecasts based on a subset of the information available to them, yet behave as if these forecasts were based on complete information.

While this general approach is helpful in understanding a number of phenomena, it also has an important limitation, since it typically takes as *exogenous and unchanging* the subset of available information that an agent restricts herself to. Consider Hirshleifer and Teoh (2002), who assume that investors have limited attention and thereby focus exclusively

¹ For textbook discussions, see, e.g., Nisbett and Ross (1980) and Fiske and Taylor (1991). We review this and related work in more detail below.

² A partial list includes: i) Miller (1977), Harrison and Kreps (1978), Varian (1989), Kandel and Pearson (1995), Morris (1996), Odean (1998), Kyle and Wang (1997), Hong and Stein (2003), and Scheinkman and Xiong (2003), all of whom couch their models in terms of either differences of opinion or overconfidence; ii) Hong and Stein (1999), who appeal to bounded rationality; and iii) Hirshleifer and Teoh (2002) and Sims (2003), who invoke limited attention.

on a firm's reported earnings, while ignoring other numbers and footnotes. One way to interpret this assumption is to think of investors believing in what we might call a "Truthful-Earnings" (TE) model.³ This model's premise is that current accounting earnings are an honest reflection of what is going on inside the firm, and hence a good basis for extrapolating future cashflows. Moreover, according to the TE model, other pieces of information, such as the structure of the CEO's compensation package, are irrelevant—the board of directors is assumed to be setting CEO compensation optimally, and whatever the consequences of such compensation for profitability, they must already be impounded in earnings.

If investors make forecasts using the TE model, it follows that managers will have incentives to manipulate these forecasts, by using various accounting tricks to artificially inflate reported earnings. This is the point that Hirshleifer and Teoh (2002) emphasize. But such a conclusion raises a series of questions. Won't even relatively naïve investors eventually learn that their simple TE model is flawed, say, following the highly-publicized scandals at Enron, WorldCom, Tyco, etc.? If so, how will they change the model they use over time? Will they eventually become more cynical, and adopt a new model that pays less attention to reported earnings, and more attention to numbers that may help flag accounting manipulation or other forms of misbehavior?

Our goal in this paper is to begin to address these kinds of questions. As in previous work, we start with the assumption that agents use simple models, i.e., models that consider only a subset of available information. But unlike this other work, we then go on to explicitly analyze the process of learning and model change. In particular, we assume that agents keep

³ We should emphasize that this is our spin, not theirs. Hirshleifer and Teoh (2002) argue that investors restrict their attention to items that are particularly salient, so that it is the salience and visibility of earnings—as opposed to the less visible kinds of numbers reported in footnotes—that matters.

track of the forecast errors associated with their simple models. If a given model performs poorly over a period of time, it may be discarded in favor of an alternative model—albeit an equally oversimplified one—that would have done better over the same period.

To be more precise, our set-up can be described as follows. Imagine a stock that at each date t pays a dividend of $D_t = A_t + B_t + \varepsilon_t$, where A_t and B_t can be thought of as two distinct sources of public information, and where ε_t is random noise. The idea that an agent uses an oversimplified model of the world can be captured by assuming that her forecasts are based on either the premise that: i) $D_t = \gamma A_t + \varepsilon_t$ (we call this having an “A model”); or ii) $D_t = \gamma B_t + \varepsilon_t$ (we call this having a “B model”). Suppose the agent initially starts out with the A model, and thus focuses only on information about A_t in generating her forecasts of D_t . Over time, the agent keeps track of the forecast errors that she incurs with the A model, and compares them to the errors she would have made had she used the B model instead. Eventually, if the A model performs poorly enough relative to the B model, we assume that the agent switches over to the B model; we term such a switch a “paradigm shift”.⁴

This type of learning is Bayesian in spirit, and we use much of the standard Bayesian apparatus to formalize the learning process. However, there is a critical sense in which our agents are not conventional fully rational Bayesians: we allow them to update only over the class of simple univariate models. That is, their priors assign zero probability to the correct

⁴ Our rendition of the learning process is inspired in part by Thomas Kuhn’s (1962) classic, The Structure of Scientific Revolutions. Kuhn argues that scientific observation and reasoning is shaped by simplified models, which he refers to as paradigms. During the course of what Kuhn calls “normal science”, a single generally-accepted paradigm is used to organize data collection and make predictions. Occasionally, however, a crisis emerges in a particular field, when it becomes clear that there are significant anomalies that cannot be rationalized within the context of the existing paradigm. According to Kuhn, such crises are ultimately resolved by revolutions, or changes of paradigm, in which an old model is discarded in favor of a new one that appears to provide a better fit to the data.

multivariate model of the world, so no matter how much data they see, they can never learn the true model.⁵

This assumption yields a range of empirical implications, which, for the sake of concreteness, we develop in a stock-market setting. Indeed, even before introducing learning effects, the assumption that agents use oversimplified models allows us to parsimoniously capture some of the best-known patterns in stock returns, such as momentum (Jegadeesh and Titman (1993)), and the value/glamour differential, or book-to-market effect (Fama and French (1992), Lakonishok, Shleifer and Vishny (1994)).

Nevertheless, the primary contribution of the paper lies in delineating the additional effects that arise from our learning mechanism. We highlight five of these. First, the book-to-market effect is amplified. Second, there is substantial variation in the conditional intensity of the book-to-market effect. For example, when a high-priced glamour stock has recently experienced a string of negative earnings surprises, there is an increased probability of a paradigm shift that will tend to be accompanied by a large negative return. Thus the conditional expected return on the stock is more strongly negative than would be anticipated on the basis of its high price alone.

The same reasoning also yields our third and fourth implications—that, even with symmetric and homoskedastic fundamentals, both the volatility and skewness of returns are stochastic, with movements that can be partially forecasted based on observables. In the above example of a glamour stock that has experienced a series of negative earnings shocks,

⁵ The idea that agents attempt to learn, but assign zero probability to the true model of the world, is also in Barberis, Shleifer and Vishny (1998). We discuss the connection between our work and this paper below.

the increased likelihood of a paradigm shift corresponds to elevated conditional volatility as well as to negative conditional skewness.

And finally, these episodes will be associated with a kind of *revisionism*: when there are paradigm shifts, investors will tend to look back at old, previously-available public information, and to draw very different inferences from it than they had before. In other words, when asked to explain a dramatic movement in a company's stock price, observers may point to data that has long been in plain view in the company's annual reports, but that was overlooked under the previous paradigm.

Several of our predictions run closely parallel to standard accounts of stock-price bubbles. In particular, the idea that a high-priced glamour stock is especially vulnerable to a large downward correction (and hence has negatively skewed returns) is also a feature of Blanchard and Watson's (1982) work on stochastic bubbles. But whereas the popping of the bubble is exogenous in their framework, our theory endogenizes it. Moreover, in so doing, we are able to say more about the exact circumstances in which a bubble will pop—i.e., after a string of bad earnings shocks that call into question the current valuation model.

In developing our stock-market results, we consider two polar cases regarding the degree of heterogeneity among investors. At one extreme, we examine a setting where there is a single representative agent, so the market price is just given by this agent's valuation. In this case, we do not allow the agent to make blended forecasts using a weighted combination of the A and B models; if she did, her forecasts would no longer satisfy our simplicity criterion of making use of just one source of information. Rather, our representative agent does the same thing that researchers in economics and many other scientific fields typically do

when they need to make model-based forecasts: she engages in *model selection*—i.e., picking a single favored model—as opposed to *Bayesian model averaging*.⁶

The representative-agent case is helpful in drawing out the intuition behind our results, so we go through it in some detail. But this approach naturally raises the question of how well our conclusions stand up when there is heterogeneity across investors. Therefore, we also consider a case in which there is a continuum of investors, each of whom has a different threshold for switching from one model to another. In this case, even though each individual investor still practices model selection, the market as a whole effectively practices a form of Bayesian model averaging. And interestingly, the qualitative predictions that emerge are very similar to those in the representative-agent case. This suggests that the key to these results is not the distinction between model selection vs. model averaging, but rather the fact that, in either case, we restrict the updating process to the space of simple models.

The rest of the paper is organized as follows. Section II briefly reviews some of the literature in psychology that is most relevant for our purposes. In Section III, we lay out our theory, and develop its implications for a variety of stock-return patterns. In Section IV, we present a case study of the recent history of Amazon.com, which provides an illustration of the basic paradigm-shift mechanism underlying our theory, as well as of the associated phenomenon of revisionism. Section V looks at the connection between our work and several related papers, and Section VI concludes.

⁶ Or said differently, the representative agent practices a specific form of the categorical thinking described by Mullainathan (2000): after weighing the evidence for the A and B models, she decides that one model is “right”, and one is “wrong”, and then proceeds to use the “right” model exclusively.

II. Some Evidence From Psychology

A. On the use of simple models

The idea that people use overly simplified models of the world is a fundamental one in the field of social cognition. According to the “cognitive miser” view, which has its roots in the work of Simon (1982), Bruner (1957), and Kahneman and Tversky (1973), humans are seen as having to confront an infinitely complex and ever-changing environment, endowed with a limited amount of processing capacity. In order to conserve on scarce cognitive resources, they use theories, or schema, to organize the data and make predictions.

Schank and Abelson (1977), Abelson (1978), and Taylor and Crocker (1980) review and classify these knowledge structures, and highlight some of their strengths and weaknesses. These authors argue that theory-driven/schematic reasoning helps people to do better at a number of tasks, including: the interpretation of new information; storage of information in memory and subsequent retrieval; the filling-in of gaps due to missing information; and overall speed of processing. At the same time, there are also several disadvantages, such as: incorrect inferences (due, e.g. to stereotyping); oversimplification; a tendency to discount disconfirming evidence; and incorrect memory retrieval.⁷

⁷ Kuhn (1962) discusses an experiment by Bruner and Postman (1949) in which individual subjects are shown to be extremely dependent on *a priori* models when encoding the most simple kinds of data. In particular, while subjects can reliably identify standard playing cards (such as a *black* six of spades) after these cards have been displayed for just an instant, they have great difficulty in identifying anomalous cards (such as a *red* six of spades) even when they are given an order of magnitude more time to do so. However, once they are aware of the existence of the anomalous cards—i.e., once their model of the world is changed—subjects can identify them as easily as the standard cards.

Fiske and Taylor (1991, p. 13) summarize the cognitive miser view as follows:

“The idea is that people are limited in their capacity to process information, so they take shortcuts whenever they can...People adopt strategies that simplify complex problems; the strategies may not be normatively correct or produce normatively correct answers, but they emphasize efficiency.”

Indeed, much of the psychology literature takes it more or less for granted that people will not use all available information in making their forecasts, and instead focuses on the specific biases that shape *which kinds* of information are most likely to be attended to. To take just one example, according to the well-known availability heuristic (Tversky and Kahneman (1973)), people tend to overweight information that is easily available in their memories—i.e., information that is especially salient or vivid.

Our theory relies on the general notion that agents disregard some relevant information when making forecasts. But importantly, it does not invoke an exogenous bias against any one type of information. Thus in our setting, A_t and B_t can be thought of as two sources of public information that are *a priori* equally salient. It is only once an agent endogenously opts to use the A model that A_t can be said to become more “available”.

B. Resistance to model change

Another prominent theme in the work on theories and schemas is that of theory maintenance. Simply put, people tend to resist changing their models, even in the face of evidence that, from a normative point of view, would appear to be strongly contradictory of these models. Rabin and Schrag (1999) provide an overview of much of this work, including the classic contribution of Lord, Ross and Lepper (1979). Nevertheless, even if people are stubborn about changing models, one probably does not want to take the extreme position that they never learn from the data. As Nisbett and Ross (1980, p. 189) write:

“Children do eventually renounce their faith in Santa Claus; once popular political leaders do fall into disfavor...Even scientists sometimes change their views....No one, certainly not the authors, would argue that new evidence or attacks on old evidence can never produce change. Our contention has simply been that generally there will be less change than would be demanded by logical or normative standards or that changes will occur more slowly than would result from an unbiased view of the accumulated evidence.”

Our efforts below can be seen as very much in the spirit of this quote. That is, while we allow for the possibility that it might take a relatively large amount of data to get an agent to change models, our whole premise is that, eventually, enough disconfirming evidence will lead to the abandonment of a given model, and to the adoption of a new one.

While the idea of theory maintenance is well-developed, the psychology literature seems to have produced less of a consensus as to when and how theories ultimately change. Lacking such an empirical foundation, our approach here is intended to be as axiomatically neutral as possible. We measure the accumulated evidence against a particular model like a Bayesian would, as the updated probability (given the data and a set of priors) that the model is wrong. And when this probability reaches a pre-determined critical value, we assume that the model is discarded. In keeping with the principle of theory maintenance, this critical value can be set to be very close to one, in which case a model is rejected only when it appears to be almost certainly wrong. But we do not impose any further biases in terms of which sorts of data get weighted more or less heavily in the course of the Bayesian-like updating.

III. Theory

A. Basic Ingredients

We consider a single traded asset, which might represent either an individual stock, or the market as a whole. There is an infinite horizon, and at each date t , the asset pays a dividend of $D_t = F_t + \varepsilon_t \equiv A_t + B_t + \varepsilon_t$, where A_t and B_t can be thought of as two distinct

sources of public information, and where ε_t is random noise. Each of the sources of information follows an AR1 process, so that $A_t = \rho A_{t-1} + a_t$, and $B_t = \rho B_{t-1} + b_t$, with $\rho < 1$. The random variables a_t , b_t , and ε_t are all independently normally distributed, with variances of v_a , v_b , and v_ε , respectively. For the sake of symmetry and simplicity, we restrict ourselves to the case where $v_a = v_b$ in what follows.

Immediately after the dividend is paid at time t , investors see the realizations of a_{t+1} and b_{t+1} , which they can use to estimate the next dividend, D_{t+1} . Assuming a constant discount rate of r , this dividend forecast can then be mapped directly into an ex-dividend present value of the stock at time t . For a fully rational investor who understands the true structure of the dividend process, and who uses both sources of information, the ex-dividend value of the stock at time t , which we denote by V_t^R , is given by: $V_t^R = k(A_{t+1} + B_{t+1})$, where $k = 1/(1+r-\rho)$ is a dividend-capitalization multiple.

By contrast, we assume that investors use overly simplified univariate models to forecast future dividends, and hence to value the stock. In particular, at any point in time, any individual investor bases her forecast on one of two premises: i) the dividend process is $D_t = \gamma A_t + \varepsilon_t$ (we call this having an ‘‘A model’’); or ii) the dividend process is $D_t = \gamma B_t + \varepsilon_t$ (we call this having a ‘‘B model’’). Thus an investor using the A model at time t has an ex-dividend valuation of the stock, V_t^A , which satisfies $V_t^A = k\gamma A_{t+1}$, and an investor using the B model at time t has a valuation V_t^B , where $V_t^B = k\gamma B_{t+1}$.⁸

⁸ Note that another possible univariate model is to forecast future dividends based solely on observed values of past dividends. That is, one can imagine a ‘‘D model’’ where $V_t^D = kD_t$. As a normative matter, the D model may be more accurate than either the A or the B model. (This happens when v_ε is small relative to the variances of A_t and B_t .) But given their mistaken beliefs about the structure of the dividend process, agents will always consider the D model to be dominated by both the A and the B models.

In what follows, we consider values of the parameter γ that satisfy $\gamma \geq 1$, although much of our focus is on the limiting case where $\gamma = 1$. On the one hand, one might claim that $\gamma = 1$ is a reasonable point of departure, because it implies that while ignoring one source of information, investors at least put the appropriate weight on the source that they do consider. Alternatively, one might argue that if an investor is only going to consider a small subset of available information to be useful, it might be quite natural for her to exaggerate the importance of that information, which would correspond to $\gamma > 1$.⁹ However, as it turns out, most of the qualitative results that we are interested in do not depend on the exact value of γ . When we do restrict attention to $\gamma = 1$, we do so only because it makes the intuition more transparent without changing any of the general conclusions.

We will soon have much more to say about the learning process which pins down the model—either A or B—that an investor uses at any given point in time. However, as a benchmark, we begin by looking at how things work when each investor uses an exogenously specified model that never changes.

B. Benchmark Case: No Learning

The simplest no-learning situation is one in which there is a single representative investor. Without loss of generality, we can assume that this investor always uses the A

⁹ In particular, another plausible benchmark is $\gamma = \sqrt{2}$, which has the property that investors' forecasts have the same variance as fully rational forecasts. This case is intuitively appealing if investors know v_ε and can get a sense of the overall variance of dividends relatively easily. For if they were to use a model with a value of γ significantly less than $\sqrt{2}$, such a model would produce forecasts that appeared too smooth relative to the observed variability in dividends.

model, so that the stock price at time t , P_t , is given by $P_t = V_t^A = k\gamma A_{t+1}$. The excess return from $t-1$ to t , which we denote by R_t , is defined by $R_t = D_t + P_t - (1+r)P_{t-1}$. It is straightforward to show that we can rewrite R_t as $R_t = z_t^A + k\gamma a_{t+1}$, where z_t^A is the forecast error associated with trying to predict the time- t dividend using model A, i.e., where $z_t^A = B_t - (\gamma-1)A_t + \varepsilon_t$. That is, under the A model, the excess return at time t has two components: i) the forecast error z_t^A ; and ii) the incremental A-news about future dividends, $k\gamma a_{t+1}$.

With these variables in hand, various properties of stock returns can be immediately established. Consider first the autocovariance of returns at times t and $t-1$. We have that:

$cov(R_t, R_{t-1}) = cov(z_t^A, z_{t-1}^A) + k\gamma cov(z_t^A, a_t)$. With a little manipulation, this yields:

$$cov(R_t, R_{t-1}) = \rho(v_b + v_a(\gamma-1)^2)/(1 - \rho^2) - kv_a\gamma(\gamma-1) \quad (1)$$

The first term in (1), $\rho(v_b + v_a(\gamma-1)^2)/(1 - \rho^2)$, is always positive, and reflects a “repeating-the-same-mistake” effect. Since the investor uses the same wrong model to make forecasts for times $t-1$ and t , her forecast errors, z_{t-1}^A and z_t^A , will be positively correlated, which tends to induce positive autocovariance in returns. The second term, $-kv_a\gamma(\gamma-1)$, is always negative, and reflects an overreaction effect. To the extent that the investor puts too much weight on A information in forecasting future dividends (i.e., $\gamma > 1$) there will tend to be reversals when actual dividends are realized.

In the polar case where $\gamma = 1$, the latter overreaction effect disappears, and (1) simplifies to $cov(R_t, R_{t-1}) = \rho v_b/(1 - \rho^2)$. Now all that is left is the positive autocovariance associated with making the same mistake—i.e., ignoring the persistent B information—in

every period. More generally, for $\gamma > I$, (and assuming that $v_b = v_a$), the autocovariance will be positive if $\rho(I + (\gamma-1)^2)/(1 - \rho^2) > k\gamma(\gamma-1)$, and negative otherwise.

The analysis of autocovariances at more distant lags is very similar. In fact, it is easy to show that, for any $j > I$, we have:

$$\text{cov}(R_b, R_{t-j}) = \rho^{j-1} \{ \rho(v_b + v_a(\gamma-1)^2)/(1 - \rho^2) - kv_a\gamma(\gamma-1) \} = \rho^{j-1} \text{cov}(R_b, R_{t-1}) \quad (2)$$

In other words, autocovariances at all horizons have the same sign, and their magnitude decays smoothly with the lag length j .

Another item of interest is the covariance between the price level and future returns, i.e., $\text{cov}(R_b, P_{t-1})$. Since all dividends are paid out immediately as realized (there are no retained earnings), and since the scale of the dividend process never changes over time, it makes sense to think of the stock as a claim on an asset with a constant underlying book value. Thus one can interpret the price of the stock—which is stationary in our model—as an analog to the market-to-book ratio, and $\text{cov}(R_b, P_{t-1})$ as a measure of how strongly this ratio forecasts returns. We can show that:

$$\text{cov}(R_b, P_{t-1}) = -kv_a\gamma(\gamma-1)/(1 - \rho^2) \quad (3)$$

As long as $\gamma > I$, the covariance between the price level and future returns is negative, implying the familiar value/glamour differential. The intuition is again one of simple overreaction: when $\gamma > I$, the price puts too much weight on A information, a mistake that will have to be gradually reversed as actual dividends are realized.

We take the following message away from the analysis to this point. Even before adding any learning considerations, we are able to capture two of the most fundamental patterns that have been documented in stock prices, just by invoking the assumption that agents use overly simplified models. So perhaps our formulation can be said to be on the right track.¹⁰ But several other recent behavioral theories can also jointly explain momentum and the value/glamour differential, and our ultimate aim is to speak to a broader set of phenomena.¹¹ Thus we would like to focus the reader's attention on the additional mileage that we get once we introduce our form of learning; what we have so far should be seen only as a (hopefully sensible) point of departure.

Indeed, in much of our analysis of the case with learning below, we simplify things by setting $\gamma = 1$. This parameterization makes the no-learning case less interesting and realistic in its own right, but at the same time allows for a cleaner and easier-to-interpret benchmark. For example, with $\gamma = 1$, the no-learning case never generates a value/glamour differential, so if we see any such differential in what follows, it will be clear that it is entirely the product of the dynamics that result from our form of learning.

¹⁰ The no-learning case can be enriched by allowing for heterogeneity among investors. Suppose a fraction f of the population use Model A, and $(1 - f)$ use model B. We can demonstrate that this set-up still generates momentum in stock returns. More interestingly, momentum is strongest when there is maximal heterogeneity among investors, i.e. when $f = 1/2$. Since such heterogeneity also generates trading volume, we have the prediction that momentum will be greater when there is more trading volume, which fits nicely with the empirical findings of Lee and Swaminathan (2000). Although this extension of the no-learning case strikes us as promising, we do not pursue it in detail here, as our main goal is to draw out the implications of our particular learning mechanism.

¹¹ See, e.g., Daniel, Hirshleifer and Subrahmanyam (1998), Barberis, Shleifer and Vishny (1998) and Hong and Stein (1999).

C. Learning: Further Ingredients

To introduce learning, we must specify several further assumptions. The first of these is that at any point in time t , an agent believes that the dividend process is governed by either the A model or the B model—i.e., she believes that either $D_t = \gamma A_t + \varepsilon_t$, or that $D_t = \gamma B_t + \varepsilon_t$. The crucial point is that the agent always wrongly thinks the true process is a univariate one, and attaches zero probability to the correct, bivariate model of the world.

Second, the agent believes that the underlying dividend process switches over time—between being driven by the A model vs. the B model—according to a Markov chain. Let π_A be the conditional probability that the agent attaches to dividends being generated by the A model in the next period given that they are being generated by the A model in the current period, and define π_B symmetrically. To keep things simple, we set $\pi_A = \pi_B = \pi$, and assume that $\frac{1}{2} < \pi < 1$, which means that the agent thinks that both states are persistent but not perfectly absorbing.

The latter assumption is not really necessary for our principal results; we can alternatively work with the limiting case of $\pi = 1$, in which the agent (correctly) thinks that nature is unchanging—i.e., that there is only a single model that applies for all time. As will become clear, the only advantage of keeping $\pi < 1$ is that it makes the learning process stationary and thereby gives our results a more steady-state flavor. In particular, with $\pi < 1$, the probability of a paradigm shift will not be a function of how many periods have elapsed since the learning process started. By contrast, with $\pi = 1$, learning is non-stationary: after a long stretch of time, there is a high probability that the agent will be almost convinced by one of the two models, thereby making further paradigm shifts extremely unlikely.

With these assumptions in place, a first step is to describe how Bayesian updating works, given the structure and the set of priors that we have specified. It is important to stress that in our setting, one does not want to interpret such Bayesian updating as corresponding to the behavior of a fully rational agent, since we have restricted the priors in such a way that no weight can ever be attached to the correct model of the world. Let p_t be the probability weight on the A model going into period t . To calculate the posterior after period t , recall that for each model, we can construct an associated forecast error, with $z_t^A = B_t - (\gamma-1)A_t + \varepsilon_t$ being the error from the A model, and $z_t^B = A_t - (\gamma-1)B_t + \varepsilon_t$ being the error from the B model. Intuitively, the updating process should tilt more in the direction of model A after period t if z_t^A is smaller than z_t^B in absolute value, and vice-versa.

Define $x_{t+1} = p_t L_z / (p_t L_z + (1-p_t))$, where L_z is the likelihood function given by $L_z = \exp(- [(z_t^A)^2 - (z_t^B)^2] / 2v_\varepsilon)$. Standard techniques can be used to calculate the Bayesian posterior going into period $t+1$:

$$p_{t+1} = p^* + (\pi_A + \pi_B - 1)(x_{t+1} - p^*) \quad (4)$$

where $p^* = (1 - \pi_B) / (2 - \pi_A - \pi_B)$ is the fraction of the time that the dividend process is expected to spend in the A-model state over the long run. Given our assumption that $\pi_A = \pi_B$, it follows that $p^* = 1/2$, and (4) reduces to:

$$p_{t+1} = 1/2 + (2\pi - 1)(x_{t+1} - 1/2) \quad (5)$$

Observe that in the limiting case where $\pi = 1$, we have that $p_{t+1} = x_{t+1}$. This is the point mentioned earlier—that Bayesian beliefs in this case are non-stationary, and eventually drift towards a value of either zero or one. However, as long as $\pi < 1$, Bayesian beliefs are stationary, with a long-run mean weight of $\frac{1}{2}$ being attached to the A model. In either case, however, the central intuition to retain is that the updating process leans more towards the A model after period t if z_t^A is smaller than z_t^B in absolute value, and vice-versa.

D. Representative-Investor Case: The Market as Model Selector

As noted above, we assume that any individual agent practices model selection. Thus if the market as a whole can be thought of in terms of a single representative investor, the price at any point in time will be given by the representative investor's valuation, according to whichever model she is currently using. To capture this idea, we stipulate that at time t , the representative investor has a preferred null model, which she uses exclusively. Moreover, as long as the accumulated evidence against the null model is not too strong, it is carried over to time $t+1$.

To be more precise, we define the indicator variable I_t^A to be equal to one if the investor's null model at time t is the A model, and to be equal to zero if it is the B model. We then assume the following dynamics for I_t^A :

$$\text{If } I_t^A = 1, \text{ then } I_{t+1}^A = 1, \text{ unless } p_{t+1} < h \quad (6)$$

$$\text{If } I_t^A = 0, \text{ then } I_{t+1}^A = 0, \text{ unless } p_{t+1} > (1 - h) \quad (7)$$

Here h is a critical value that is less than one-half. Thus the investor maintains a given null model for the purposes of making forecasts until the updated (Bayesian) probability of it being correct falls below the critical value. So, for example, if her original null is the A model, and $h = 0.05$, she continues to make forecasts exclusively with it until it is rejected at the five-percent confidence level. Once this happens, the B model assumes the status of the null model, and it is then used exclusively until it too is rejected at the five-percent confidence level. Clearly, the smaller is h , the stronger is the degree of resistance to model change. So one way to rationalize a value of h close to zero is by appealing to the psychological literature on theory maintenance discussed above.

This formulation raises an important issue of interpretation that we have thus far glossed over. On the one hand, we have tried to motivate the assumption that the investor uses a univariate forecasting model at any point in time by appealing to limited cognitive resources—the notion being that it is too difficult to simultaneously process both the A and B sources of information for the purposes of making a forecast. Yet at the same time, the investor does use both the A and B sources of information when deciding whether to abandon her null model—the Bayesian updating process for p_t which underlies her model-selection criterion depends on both z_t^A and z_t^B . In other words, the investor is capable of doing quite sophisticated multivariate operations when evaluating which model is better, but is unable to make dividend forecasts based on more than a single variable at a time, which all sounds somewhat schizophrenic.

One resolution to this apparent paradox relies on the observation that, in spite of the way we have formalized things, it is neither realistic nor necessary for our results to have the representative investor actively review her choice of models as frequently as once every

period. Indeed, it is more plausible to think of the two basic tasks that the investor undertakes—forecasting and model selection—as happening on very different time scales, and therefore involving fundamentally different tradeoffs of cognitive costs and benefits. For an active stock-market participant, dividend forecasts have to be updated continuously, as new information comes in. Thus the model that generates these forecasts needs to be simple and not too cognitively burdensome, or it will be impractical to use it in real time.¹²

In contrast, it may well be that the investor steps back from the ongoing task of forecasting and does systematic model evaluation only once in a long while; as a result, it might be feasible for this process to be more data-intensive.¹³ Indeed, it is not difficult to incorporate this sort of timing feature explicitly into our analysis, e.g., by allowing the investor to engage in model evaluation only once every m periods, with m relatively large. Our limited efforts at experimentation suggest that this approach yields results that are qualitatively similar to those we report below.

E. Heterogeneous-Investor Case: The Market as Model Averager

As will become clear, the representative-investor/model-selection approach described above provides a useful way to communicate the main intuition behind our results. But it is important to underscore that these results do not hinge on the discreteness associated with the

¹² This is why we are reluctant to assume that any individual agent acts as a model averager. If a model averager assigns a probability p_t to the A model at time t , her forecast of the next dividend would be $p_t \mathcal{A}_{t+1} + (1 - p_t) \mathcal{B}_{t+1}$. However, such a forecast is no longer a cognitively simple one to make in real time, as it requires the agent to make use of both sources of information simultaneously. And if we are going to endow the agent with this much high-frequency processing power, it is no longer clear how one motivates the assumption that she does not consider more complicated models in her set of priors.

¹³ Moreover, much of this low-frequency model evaluation may happen at the level of an entire investment community, rather than at the level of any single investor. For example, each investor may need to work alone with a given simple model to generate her own high-frequency forecasts, but may once in a while change models based on what she reads in the press, hears from fellow investors, etc. Again, the point to be made is that no single investor is literally going to be engaging in cognitively costly model evaluation on a continuous basis.

model-selection mechanism. To illustrate this point, we also consider the “smoother” case where the market price is based on model averaging, i.e., where $P_t = p_t k \gamma A_{t+1} + (1 - p_t) k \gamma B_{t+1}$. This model-averaging case can be motivated by appealing to a particular form of heterogeneity across investors.

To see this, suppose that there are a continuum of investors distributed uniformly across the interval $[0, 1]$, each of whom individually practices model selection. All investors share the same underlying Bayesian update p_t of the probability of the A model being correct at time t , with p_t evolving as before. But now, each investor has her own fixed threshold for determining when to use the A model as opposed to the B model: the investor located at point i on the interval uses the A model if and only if $p_t > i$.¹⁴ This implies that the fraction of investors in the population using the A model at time t is given by p_t . And to the extent that the market price is just the weighted average of individual investors’ estimates of fundamental value, this in turn implies that $P_t = p_t k \gamma A_{t+1} + (1 - p_t) k \gamma B_{t+1}$.¹⁵

F. Implications of Learning for Stock Returns

Unlike in the no-learning case, we are no longer able to solve for various moments of interest in closed form, and we have to resort to computer simulations to approximate these moments for any given set of parameters. Nevertheless, it is possible to draw out the intuition for our results in some detail. We do this—largely in the context of the representative-investor case—before proceeding to the numerical examples.

¹⁴ One can interpret investors with low thresholds as those who have an innate preference for the A model.

¹⁵ This motivation is admittedly loose. In a dynamic model, it is not generally true that price simply equals the weighted average estimate of fundamental value—short-term-trading considerations arise, as, e.g., investors try to forecast the forecasts of others. Nevertheless, since we just want to demonstrate that our results are not wholly dependent on model selection, the simple model-averaging case is a natural point of comparison.

1. *Representative-Investor/Model-Selection Case*

The intuition is most transparent when we set $\gamma = 1$. Suppose for the moment that the representative investor is using the A model at time $t-1$, so that $P_{t-1} = kA_t$. There are two possibilities at time t . The first is that there will be no paradigm shift, so that the investor continues to use the A model. In this case, $P_t = kA_{t+1}$, and the return at time t , which we denote by R_t^N , is given by:

$$R_t^N = z_t^A + ka_{t+1} = B_t + \varepsilon_t + ka_{t+1} \quad (8)$$

Alternatively, if there is a paradigm shift at time t , the investor switches over to using the B model, in which case the price is $P_t = kB_{t+1}$, and the return, denoted by R_t^S , is:

$$R_t^S = z_t^A + kb_{t+1} + \rho k(B_t - A_t) = B_t + \varepsilon_t + kb_{t+1} + \rho k(B_t - A_t) \quad (9)$$

Observe that $R_t^S = R_t^N + k(b_{t+1} - a_{t+1}) + \rho k(B_t - A_t)$. Simply put, the return in the paradigm-shift case differs from that in the no-shift case as a result of current and lagged A-information being discarded from the price, and replaced with B-information.

a. *The value/glamour differential*

Let us begin by revisiting the magnitude of the value/glamour effect, as proxied for by $cov(R_t, P_{t-1})$. (Recall that for $\gamma = 1$, we had $cov(R_t, P_{t-1}) = 0$ in the no-learning case.) We can decompose $cov(R_t, P_{t-1})$ as follows:

$$\begin{aligned} cov(R_t, P_{t-1}) &= cov(R_t^S, P_{t-1}/shift) * prob(shift) + \\ &cov(R_t^N, P_{t-1}/no\ shift) * prob(no\ shift) \end{aligned} \quad (10)$$

Substituting in the definitions of R_t^S and R_t^N from (8) and (9), and simplifying, we can rewrite (10) as:

$$\begin{aligned} cov(R_t, P_{t-1}) &= k \{ cov(\varepsilon_t, A_t) + cov(A_t, B_t) \} + \\ &\rho k^2 \{ cov(A_t, B_t/shift) - var(A_t/shift) \} * prob(shift) \end{aligned} \quad (11)$$

Note that both the $cov(\varepsilon_t, A_t)$ term, as well as the first $cov(A_t, B_t)$ term in (11), are *unconditional* covariances. We have been assuming all along that these unconditional covariances are zero. Thus (11) can be further reduced to:

$$cov(R_t, P_{t-1}) = \rho k^2 \{ cov(A_t, B_t/shift) - var(A_t/shift) \} * prob(shift) \quad (11')$$

Equation (11') clarifies the way in which a value/glamour effect arises when there is learning. A preliminary observation is that $cov(R_t, P_{t-1})$ can only ever be non-zero to the extent that the probability of a paradigm shift, $prob(shift)$, is non-zero: as we have already mentioned, when $\gamma = 1$, there is no value/glamour effect absent learning. When $prob(shift) > 0$, there are two distinct mechanisms at work. First, there is the negative contribution from the $-var(A_t/shift)$ term. This term reflects the fact that A-information is abruptly removed from the price at the time of a paradigm shift. This tends to induce a negative covariance between the price level and future returns, since, e.g., a highly positive value of A_t at time $t-1$ will lead

to a high price at this time, and then to a large negative return when this information is discarded from the price at time t .

Second, and more subtly, there is the $cov(A_t, B_t/\text{shift})$ term. Of course, the unconditional covariance between A_t and B_t is zero. But the covariance conditional on a paradigm shift is not. To see this heuristically, think about the circumstances in which a shift from the A model to the B model is most likely to occur. Such a shift will tend to happen after a series of realizations for which z_t^A has been significantly larger in absolute value than z_t^B . Now recall that for $\gamma = 1$, we have $z_t^A = B_t + \varepsilon_t$, and $z_t^B = A_t + \varepsilon_t$. So clearly, when A_t and B_t are very close together (think of the limiting case where $A_t = B_t$), there is little scope for a paradigm shift, no matter what the realization of ε_t . Said differently, it is hard to learn that one model is better than the other if both models are making the same forecast.

In contrast, if the two models make opposing forecasts, there is room for learning. Moreover, if luck favors the B model, this learning will tend to induce a shift away from the A model. As a simple example, consider a fixed value of A_t , and imagine that we have both $B_t = -A_t$, and $\varepsilon_t = -A_t$: this makes $z_t^A = -2A_t$, and $z_t^B = 0$, and raises the likelihood that the A model will be discarded.

This line of reasoning suggests that $cov(A_t, B_t/\text{shift}) < 0$, which makes the overall value of $cov(R_t, P_{t-1})$ in (11') even more negative—i.e., it strengthens the value/glamour differential. When a paradigm shift occurs, not only is A-information discarded from the price, it is also replaced with B-information. And conditional on a shift occurring, these two pieces of information tend to be pointing in opposite directions. So if a positive value of A_t at $t-1$ has led to a high price at this time, there will tend to be an extra negative impact on returns in the

event of a paradigm shift at t —above and beyond that associated with just the discarding of A_t —when B_t enters into the price for the first time.

b. *Conditional variation in the value/glamour differential*

In our setting, learning does more than just strengthen the value/glamour effect. It also introduces predictable variation in the intensity of this effect. To see why, note that much of the value/glamour effect is concentrated in those periods when paradigm shifts occur. (Indeed, the effect is entirely concentrated in such periods if $\gamma = 1$.) Thus if an econometrician can track variation over time in the probability of a paradigm shift—i.e., if he can predict when the potential for learning is relatively high—he will also be able to forecast when the value/glamour differential is likely to be greatest.

We have already seen that there is more potential for learning when the A model and the B model make divergent forecasts. What does this mean in terms of observables? To be specific, think of a situation in which A_t is very positive, so the stock is a high-priced glamour stock. Going forward, there will be more scope for learning if, in addition, B_t is negative. This will tend to show up as negative values of the forecast error z^A_t , since $z^A_t = B_t + \varepsilon_t$. In other words, if a high-priced stock is experiencing negative forecast errors, this is a clue that the two models are at odds with one another.

Thus a sharper prediction of our theory is that a high-priced glamour stock will be particularly vulnerable to a paradigm shift—and hence to a sharp decline in prices—after a series of negative z -surprises about fundamentals. Conversely, a low-priced value stock will be most likely to experience a sharp price increase after a series of positive z -surprises. The closest empirical analog to such z -surprises would probably be either: i) a measure of realized earnings in a given quarter relative to the median analyst's forecast for earnings; or ii) the

stock-price response on the day of an earnings announcement. However, as we demonstrate in our simulations below, we also get similar conclusions, albeit with slightly reduced magnitudes, if we instead define surprises in terms of recent stock returns (i.e., total returns over an interval, not just the component of returns due to an earnings announcement). This makes sense, given the connection between z -surprises and stock returns: recall that if investors are currently using the A model, $R_t = z_t^A + ka_{t+1}$, which means that the return is just a noisy version of the z -surprise.

When we say that a glamour stock has more negative expected returns conditional on a recent string of disappointing earnings surprises or disappointing past returns, we need to stress a crucial distinction. This phenomenon is *not* simply a result of adding together the unconditional value/glamour and momentum effects. Rather, in the context of a regression model to forecast future returns, our theory predicts that not only should there be book-to-market and momentum variables, but also an *interaction* of the book-to-market variable with a “momentum-like” measure, ideally one that captures the direction of recent earnings surprises. We will highlight this distinction in our simulations, by showing that the conditional variation in the value/glamour differential that we are talking about still arises for parameter configurations such that the unconditional momentum effect is zero, or even negative.

Asness (1997) produces evidence which bears directly on this prediction of our theory. Using data from 1963-1994, he performs a five-by-five sort of stocks along two dimensions: glamour/value (measured with industry-adjusted book-to-market ratios); and momentum (measured using returns over the last 12 months, excluding the most recent one.) In the most negative momentum quintile, glamour stocks (i.e., those in the lowest quintile of book-to-market) underperform moderately-priced stocks (those in the middle quintile of book-to-

market) by 77 basis points per month. In contrast, in the highest momentum quintile, the corresponding underperformance figure for glamour stocks is only 1 basis point per month.¹⁶

c. Conditional variation in volatility and skewness

The same basic mechanisms produce partially forecastable movements in stock-return volatility and skewness. As a comparison of equations (8) and (9) makes clear, volatility is inherently stochastic in our setting, because returns have more variance at times of paradigm shifts than at other times. Moreover, these movements in volatility can be partially forecasted by an econometrician, using exactly the same logic as above. For example, a high-priced glamour stock is more apt to experience a paradigm shift—which will manifest itself not only as a negative return, but also as an unusually large absolute price movement—after a sequence of negative fundamental surprises. Again, this is because such negative surprises are an indicator that the A and B models are in disagreement, which raises the potential for learning.

Analogous arguments apply for conditional skewness. First, glamour stocks will tend to have more negatively skewed returns than value stocks. This is because the very largest movements in glamour stocks—i.e., those associated with paradigm shifts—will on average be negative, and conversely for value stocks. This prediction squares well with the evidence in Chen, Hong and Stein (2001), who document that, at both the level of individual stocks and the market as a whole, the skewness of daily returns is more negative when prices are high (i.e., when book-to-market ratios are low). This feature of our theory is also reminiscent of classic accounts of bubbles: the potential for the sudden popping of a bubble in a high-priced

¹⁶ These numbers are from Table 4 of Asness (1997, p. 32). Moreover, the underperformance of glamour stocks declines monotonically across momentum quintiles: from 77 basis points to 29 to 28 to 18 to 1 as we move from the lowest momentum quintile to the highest.

glamour stock similarly generates negative conditional skewness. But whereas the popping of the bubble is exogenous in, e.g., Blanchard and Watson (1982), our theory endogenizes it.

Relatedly, we have the sharper prediction—as compared to standard bubble stories—that these general skewness effects will be more pronounced if one further conditions on recent news. So, for example, the negative skewness in a glamour stock will be strongest after it has experienced a recent string of bad news. And the positive skewness in a value stock will be greatest after a string of good news. We do not know of any specific evidence that speaks to either this prediction, or the analogous one for conditional volatility, so they may represent good opportunities for “out-of-sample” tests of our theory.

2. Heterogeneous-Investor/Model-Averaging Case

Although we will not go through the algebra of the model-averaging case, the underlying intuition is very similar to that above. In the model-selection case, the notion of effective learning at the market level is dichotomous: either there is a paradigm shift in a given period, or there is not. But this discreteness is not what is driving the results. Rather, what matters for the various asset-pricing patterns is that an econometrician can forecast when there is likely to be “a lot” of learning—i.e., he can tell when the A and B models are pointing in opposite directions.

In the model-averaging case, the amount of market-wide learning that takes place is a continuous variable, but the econometrician can still partially forecast it, for the same reason as before. In particular, when a glamour stock is observed to have a series of negative earnings surprises, this suggests that there is tension between the A and B models, which in turn means that the potential for learning is high. The implications for conditional variation in

the value/glamour differential, in volatility and in skewness all follow from this ability to anticipate variation over time in the degree of learning.

G. Simulations

We now demonstrate our basic results with a series of simulations, which are shown in Table 1. In Panel A of the table, the parameters are set as follows. The number of time periods in each run is $T = 2,000$. The variances of the shocks are set equal to $v_a = v_b = v_\varepsilon = 0.0001$. The autocorrelation coefficient ρ of the A_t and B_t processes is set equal to 0.85 . And finally, $r=0.03$, $h=0.051$, $\pi_A = \pi_B = 0.95$ and $\gamma=1$. In Panels B and C, everything else is the same, but we increase ρ to 0.90 and 0.95 respectively. As will become clear momentarily, these parameters generate values of stock-return volatility that make it sensible to think of a single period as representing one calendar quarter. For each set of parameters, we generate $N = 1,000$ different time series of stock prices, and use the averages across these series to calculate a variety of statistics. We repeat these calculations across three different cases: i) the benchmark case with no learning (column 1); ii) the representative-investor/model-selection case (column 2); and iii) the heterogeneous-investor/model-averaging case (column 3).

We begin with several unconditional moments. ShiftProb (which applies only in the model-selection case) is the number of paradigm shifts divided by T —i.e., ShiftProb is the unconditional probability of a shift occurring in a given period. Volatility is the square root of $E[(R_{t+1})^2]$; β^{MOM} is the coefficient in a regression of R_{t+1} on the cumulative return over the prior four periods, $R_{t-3,t}$; β^{VALUE} is the coefficient in a regression of R_{t+1} on P_t ; and Expected Return/Glamour is the expected excess return on a stock with an above-average price, i.e., $E[R_{t+1} | P_t > 0]$.

Next we have a series of conditional moments. In one version of these we condition both on the stock being a glamour stock and on the cumulative return due to the last four dividend announcements being negative. Thus we have Expected Return/Glamour/Bad News = $E[R_{t+1} | P_t > 0, z^*_{t-3,t} < 0]$, where $z^*_{t-3,t}$ is the cumulative return due to dividend announcements at $t-3, t-2, t-1$ and t . Similarly, Volatility/Glamour/Bad News is the square root of $E[(R_{t+1})^2 | P_t > 0, z^*_{t-3,t} < 0]$, and Skewness/Glamour/Bad News is $E[(R_{t+1})^3 | P_t > 0, z^*_{t-3,t} < 0]$. Finally, ShiftProb/Glamour/Bad News is the conditional probability of a shift in period $t+1$ given $P_t > 0$, and $z^*_{t-3,t} < 0$, and Corr(A,B)/Glamour/Bad News is the conditional correlation of A_t and B_t under the same circumstances.

In an alternative version of these moments, we condition on the stock being a glamour stock and on the cumulation of the last four *total returns*—as opposed to just the dividend-surprise components of these returns—being negative. Thus Expected Return/Glamour/Bad Returns is $E[R_{t+1} | P_t > 0, R_{t-3,t} < 0]$, with analogous redefinitions for Volatility/Glamour/Bad Returns, Skewness/Glamour/Bad Returns, ShiftProb/Glamour/Bad Returns, and Corr(A,B)/Glamour/Bad Returns.

The patterns are for the most part quite similar across the three panels A-C, so we focus our discussion on Panel B, which represents an intermediate scenario in terms of magnitudes. Consider first the case with no learning. Here the volatility of returns is 0.0809, which translates into an annualized standard deviation of 16.2% if one thinks of a period as equal to one calendar quarter. There is positive momentum in returns, with $\beta^{MOM} = 0.0505$. However, there is no value/glamour differential (i.e., $\beta^{VALUE} = 0$), an outcome which is built in by virtue of the assumption that $\gamma = 1$.

Turning to the case of model-selection-based learning, the parameters in column 2 of Panel B imply an unconditional probability of a paradigm shift of 3.92%, which corresponds to a shift approximately once every 6½ years. Volatility increases to 0.1094. The momentum effect is wiped out—in fact, it actually becomes a tiny bit negative, with $\beta^{MOM} = -0.0087$.¹⁷ On the other hand, there is now a value/glamour effect, with $\beta^{VALUE} = -0.0336$. To get a feel for the magnitude of this latter effect, note that the expected one-quarter excess return to a glamour stock is -0.0071 , or about -2.84% on an annualized basis. By symmetry, the expected return to a value stock must be the same in absolute magnitude, but with the opposite sign, implying a realistic annualized value/glamour spread of about 5.68%.

If we condition not only on a stock being a glamour stock, but also on it having had negative cumulative dividend surprises over the prior four quarters, the expected excess return goes from -0.0071 to -0.0277 , or about -11.08% on an annual basis. In other words, given the negative dividend surprises, the conditional magnitude of the glamour effect is almost four times the unconditional magnitude. Again, it should be emphasized that this result is not simply an artifact of summing the unconditional glamour and momentum effects; indeed since the unconditional momentum effect is actually slightly *negative* here, it goes the wrong way in terms of explaining the very negative expected return that obtains in this scenario.¹⁸

Under these same conditions, the probability of a paradigm shift goes up from its unconditional value of 3.92% to 5.60%, and volatility increases from 0.1094 to 0.1240.

¹⁷ This is the one place where the qualitative results are most dependent on parameter values. For example, in Panel A, there is still positive momentum even in the two cases with learning.

¹⁸ That is, given the negative value of β^{MOM} , one would all else equal expect *positive* excess returns from a stock that had recently experienced a string of bad news.

Finally, returns—which are unconditionally symmetric—become negatively skewed, with a third moment of -0.0054 .

The key to understanding these patterns is the conditional correlation of A_t and B_t . While the unconditional correlation is zero, the correlation conditional on the stock being a glamour stock with negative dividend surprises is highly negative, at -0.463 . In other words, this glamour/bad-dividend-surprise configuration is a strong signal that the A model and the B model are generating conflicting forecasts, which increases the potential for learning. This is precisely why the probability of a paradigm shift is elevated, with the accompanying implications for the other conditional moments of returns.

As an alternative to conditioning on a glamour stock having had negative cumulative dividend surprises over the prior four quarters, we also check to see what happens when we condition on it having had negative cumulative *total* returns over the same interval. One benefit of doing so is that our theoretical predictions can now be mapped directly into the evidence from Asness (1997) discussed above. The results here are similar, though slightly attenuated. For example, the expected excess return is now -0.0197 , (or -7.88% on an annual basis) as compared to the previous figure of -0.0277 . This attenuation is to be expected, because returns are a noisy proxy for dividend surprises, and it is the latter that enter directly into the investor's updating process.

In column 3, we present all the analogous results for the case of model averaging. They are for the most part remarkably similar to those for the case of model selection; indeed, if anything, the conditional patterns are a bit more pronounced. To take just one example, the expected excess return conditional on glamour and bad dividend surprises is now -0.0402 , or -16.08% per year. This makes it clear that our results are not due to the discreteness inherent

in model selection. Rather, the crucial mechanism is that when agents update over simple models, an econometrician can predict when a lot of learning is likely to take place.

IV. A Case Study: Amazon.com

In an effort to more vividly illustrate the central ideas in our theory—and to give a concrete example of the phenomenon of revisionism—we now present a case study of Amazon.com. Our focus is on the models that sell-side equity analysts have used to arrive at valuations for Amazon, and more specifically, on how these models have changed over time. All the analysts' reports that we draw on below come from the Multex database. The raw sources include reports issued by 11 different brokerage houses, at a frequency ranging from monthly to quarterly, over the period from July 1997 to December 2002.

A. Background on Amazon

Amazon, an internet retailer, was founded in 1994, opened its online store in 1995, and started trading publicly in May of 1997. Its meteoric rise and subsequent fall are well-documented. From a value of \$5 per share at year-end 1997, Amazon's stock price reached a peak of over \$106 in mid-December 1999, representing a market capitalization of roughly \$36 billion, and a multiple of over 136 times book value. (All stock-price figures are split-adjusted for comparability.) At this point, Amazon was trading for about 23 times the sum of the two largest "bricks-and-mortar" retailers of books and music, Barnes&Noble and Borders. Capping it all off, Amazon founder Jeff Bezos was named Time Magazine's "Man of the Year" for 1999.

The 1999 Christmas season—billed by some as the “first e-Christmas”—marked a turning point for Amazon. Although online shopping volume was high, Amazon failed to convert this high volume into positive profits. Its stock then went more or less straight down over the next two years, finishing the year 2001 at a price of just under \$11—about 10% of its peak value. Figure 1 plots Amazon’s stock price over the period 1997-2002.

B. Mapping Amazon Into Our Theory: What Are Models A and B?

Suppose that according to the “true” model of the world, there are two variables that, at any point in time, are useful for forecasting Amazon’s future earnings. The first variable, which we call “clicks”, is a measure of how rapidly Amazon’s customer base is growing. The second, which we call “margins”, measures how profitable incremental sales are. Intuitively, long-run profitability will by definition be given by the number of customers times the profit per customer, with clicks being a noisy predictor of the former, and margins being a noisy predictor of the latter. Of course, the precise weights that should be assigned to each of these variables will depend on a host of factors. For example, if Amazon ultimately develops a very loyal customer base and a lot of market power, future margins may exceed those earned during a period of penetration pricing, which would make current margins less informative. Nevertheless, it seems hard to argue that both the click and margin variables would not have *some* information content.

Yet if one reads the analysts’ reports, they seem to be overly fixated on a clicks-based model in the early part of Amazon’s history. These early reports dismiss the fact that Amazon’s gross margins (defined as $(\text{revenues} - \text{cost of goods sold})/\text{revenues}$) are at the time much lower than those of its closest off-line retailing peers like Barnes&Noble. In fact, they

argue repeatedly that Barnes&Noble is the wrong analogy to draw, and that Amazon should be viewed as a very different type of business.

Then, after the disappointing Christmas season of 1999, there appears to be an abrupt shift in perspective. Many analysts begin to point out the similarities between Amazon and the off-line retailers, and at the same time, start to emphasize gross margins in making their forecasts and recommendations. Indeed, a number of their post-1999 reports give a lot of play to unfavorable data on Amazon's margins that had already been widely available for some time. And strikingly, some now use this stale data to justify downgrading the stock. This is just the sort of revisionism that our theory suggests.

C. Valuation During the Bubble: A Clicks-Based Model

In a February 12, 1999 report entitled "ROIC is Key (Not the Gross Margin)" Scott Ehrens of Bear Stearns nicely summarizes the contrasting models for Amazon, and concludes that, unlike with off-line retailers, current margins are not relevant for valuation purposes.

"This is not traditional off-line retail. The gross margin is typically a good indicator of a traditional off-line retailer's return on invested capital (ROIC). However, given the highly scalable nature of the on-line model (i.e. exceptionally high revenue potential per dollar of capital invested), it becomes gross profit dollars, and not the gross margin, that drive the on-line retailer's ROIC. Amazon.com's return on invested capital, in our view, has the potential to be substantially higher than that of traditional off-line retailers....In traditional off-line retail, the gross margin is a very important metric to watch. Although a successful "bricks-and-mortar" retailer can drive sales per sq. ft higher, there is a limit to the traffic that a store can accommodate before expansion and relocation is necessary. In other words, once shelf space and traffic have been maximized, the only way to increase gross profit dollars without additional capital spending is to increase the gross margin on each sale...This is not the case in the on-line world, since an on-line retailer's revenue potential is not limited by the same factors. To illustrate, think of Amazon.com as a store with unlimited shelf space and unlimited customer base. Amazon.com does not require the same incremental capital investment to increase sales to the extent of its off-line counterpart, and, as a result, can continually increase its revenue per dollar of capital invested...This means that Amazon.com's ROIC is a function of gross profit (in absolute dollars) and not gross margin (a percentage of sales), and, therefore a higher mix of lower-margin sales could actually lead to higher ROIC given the capital efficiency of the revenue growth."

In a similar spirit is an August 4, 1998 report from Mary Meeker of Morgan Stanley Dean Witter. Meeker actually mentions that Amazon's current margins are much lower than those of Barnes&Noble. But she makes it clear that this is not relevant to her valuation, which is instead premised on a growth-oriented model similar to that which she has applied to AOL.

“Online retailing is going to be huge (already, Amazon, based on this quarter's financial results, by our math, is the second fastest growing retailer in the history of the planet), and no company is as well positioned to take advantage of the market opportunity, in part by spending to grab share (in multiple markets) early, as Amazon.com is. Translation? They are going for it. Remember, yikes, America Online spent \$1B over ten years to nab 10MM customers and it now carries a market value of \$33B...As with AOL in the early days, it's tough to determine exactly where “critical mass” is, and as long as customer addition/repeat buying/revenue generation trends remain especially positive, ongoing expense stoking is advised, because when we turn to profitability, thanks in part to economics of increasing returns, a captive customer base and scale, profit growth can be especially positive...Gross margin of 22.6% was up from 22.1% in C1Q. We continue to believe that as AMZN's buying power increases, it will be able to reach higher gross margins—the company's target is 23-27%. Remember that traditional book sellers like Barnes&Noble support gross margins near 36% due to purchasing power and, in part, due to their ability to charge higher prices in their retail locations.”

On March 9, 1999, Henry Blodgett of Merrill Lynch offers a strong recommendation of Amazon which is notable in two ways. First, like the other analysts, he explicitly rejects the analogy between Amazon and Barnes&Noble. And second, his discussion is almost entirely centered on revenue growth projections (i.e., clicks), with just an offhand nod to the assumption that net margins will eventually turn positive.

“...Amazon.com's model more closely resembles the direct sales model of computer manufacturer Dell than it does land-based retailer Barnes&Noble's....For those worried that the company will never make money, it is encouraging: Dell has an 8% net margin; Barnes&Noble...make 2%.....Our official five-year projections are similar to the Street's and assume 1) the customer base increases from 6 million to 30 million by 2003 (approximately 35% per year), 2) revenue per account increases from \$98 to \$130 by 2003 as customers buy a more diverse selection of products. Do the multiplication and—voila!—a 2003 revenue estimate of \$3.2 billion (which, when combined with an operating margin assumption of 10%, a 40X terminal multiple, and a 15% discount rate, equates to a current value of about \$30 per share.)...The risk in making conservative assumptions in this market, however, is missing a gigantic opportunity...So let's tweak those assumptions and see what happens. Let's assume that the customer base increases to 55 million and average revenue per account increases to \$170. Do that math, and suddenly, Amazon.com isn't a \$3 billion company but a \$10 billion company. Place a 12% operating margin on this revenue estimate (with additional scale, the company should be a bit more profitable), use a 50X multiple, and discount the resulting EPS back at a more aggressive 10%, and suddenly the stock is worth \$150.”

Tom Courtney of Banc of America Securities, in an August 1999 report entitled “Thinking Outside The Big Box Superstore—A White Paper on the Internet Retail Revolution”, also focuses on the growth of Amazon’s customer base as the dominant factor driving his valuation; like Blodgett, he ignores current operating margins, and simply assumes that margins will eventually turn highly positive.

“We...have identified at least one company, Amazon, that we are confident will grow at a rate that justifies its current valuation as well as significant upside potential over the next two to three years...The market is telling us—and we agree—that Internet retailers will grow at a rate far greater than the growth currently projected in most models. In fact, Internet retailing is already delivering growth that far surpasses the original expectations...We believe the growth rate will remain very strong as the number of users and buyers increases. If the online market is going to grow at 50% over the next four years, the best Internet retailers should be able to grow revenues at that rate or more. The result, based on current expectations for a number of these stocks, is that sales growth that will materially exceed the Street’s expectations. For those companies with scalable and leverageable models and strong individual transaction economics, that revenue growth will generate strong profits and ROIC.”

Finally, a rare dissenting view for this period is offered by Jonathan Cohen of Merrill Lynch, in a September 15, 1998 report. Cohen takes direct issue with the view that a retailer like Amazon will have a sticky enough customer base to justify large investments in market share at the expense of current profits. (Incidentally, soon after this report was issued, Merrill replaced Cohen with Henry Blodgett.)

“We believe that the notion that Amazon.com will be able to profitably leverage its (diminishing) market share in online book sales into other, largely unrelated business lines may prove overly optimistic...More critically, we do not believe that online commodity product sales produce the sort of brand equity generated by the distribution of proprietary information or media products. The implication here is that while it may make economic sense for Yahoo! to lose money while building a user population, it probably does not make sense for Amazon.com to follow in the same path.”

D. The Bubble Pops: the Shift to a Margins-Based Model

The Christmas season of 1999 was viewed by many as a crucial test for internet retailers. While many shoppers did go online during this season, most internet retailers,

including Amazon, made little in the way of profits. Indeed, Amazon still had negative operating income, not only for all of fiscal 1999, but also for the fourth quarter.

In the context of our theory, one might interpret the disappointing earnings results for the fourth quarter of 1999 as a low realization of D_t , one sharply at odds with the rosy forecast produced by the A (clicks-based) model, and more consistent with the forecast coming from the B (margins-based) model. According to the theory, such a configuration should be the one most likely to produce a paradigm shift in favor of the margins-based model. And consistent with this idea, analyst reports issued in early to mid-2000—including ones written by Meeker, Blodgett and Courtney—now stress that Amazon’s stock price hinges crucially on its ability to improve its margins. For example, in a report dated February 3, 2000, Genni Combes of Hambrecht and Quist writes: “Key metric going forward is gross margin improvements in domestic retail, and fulfillment costs as a percent of revenue.” And on July 27, 2000, Blodgett of Merrill Lynch lowers his rating on Amazon from a Buy to an Accumulate, saying that a key factor in this downgrade is the need for improvement in gross margin.

The striking fact here is that Amazon’s gross margins look little different in early 2000 than they did in the two years before—they had been hovering in a narrow range around 20% for quite some time.¹⁹ So there has not really been any news in the traditional sense on the gross margin front. Rather, the analysts seem to be re-evaluating the significance of previously-available information.

In part, this re-evaluation reflects a growing consensus that Amazon may not be so different from its off-line retailing peers after all. As Sara Farley of Paine Webber puts it in a February 23, 2000 report:

¹⁹ Amazon’s gross margins in 1998 varied between 22.1% and 22.6% on a quarterly basis. In 1999, the range was from 21.2% to 23.0%.

“Amazon’s customer focus has come at a price...It takes a significant amount of money, physical assets and people to provide a great shopping experience...As a result, the company’s business model has become less “virtual” over time and more physical. To be sure, Amazon still has the cost and investment advantage of not having to run physical storefronts. However, this is offset by increasing advertising spending and higher customer service costs...Putting all this together means that the company’s potential long-term operating margins are not likely to be much different than those of its off-line competitors, in the 8-10% range.”

Analysts’ relatively single-minded focus on margins continues into 2001. In a March 2, 2001 report entitled “Amazon.com: The Good, The Bad and the Ugly”, Holly Becker of Lehman Brothers writes:

“Furthermore, the company’s business model, which once promised to yield significantly higher margins than those of traditional retailers or catalogers, appears to be fundamentally disadvantaged in several areas. It is now clear that higher customer churn rates, weak shipping margins and equally high marketing spend will offset many of the company’s virtues, such as lower capital requirements and smaller labor and real estate costs. Overall, we continue to believe that Amazon’s valuation at 1.1x2000E sales and market value of \$3.7 billion remains rich, especially given the challenges facing the company....we recommend investors stay on the sidelines....Clearly, the company will need to increase gross margins to cover its fulfillment costs and make a positive contribution margin.”

Similarly, Sara D’Eathe of Thomas Weisel Partners also emphasizes gross margins in her report of April 25, 2001, which goes through an explicit valuation analysis. (Amazon is trading at a price of about \$16 per share at the time of this report.)

“Valuation is unattractive in our view. Our break-up valuation yields a price per Amazon share of \$8. We believe that Amazon’s core BMV business is worth \$3.7 billion, over 80% of its estimated market value. Assuming a discount rate of 23% and a terminal value multiple of 20x-25x P/E, our DCF model indicates a \$10-\$13 price target. Alternatively, we backed into what we believe needs to occur to gross margins in order to justify a 25% stock price return over the next 3 years. We estimate the required gross margin to be close to 39%, a level we argue is not achievable, in our view, given the merchandise mix and ongoing fulfillment inefficiencies.”

V. Related Work

There is a longstanding literature in game theory that examines the implications of learning by less-than-fully-rational agents (i.e., agents who have inconsistent and/or non-

common priors, or who may not understand the equilibria of even very simple games).²⁰ While we share some of the same behavioral premises as this work, its goals are very different than ours—for the most part, it seeks to understand the extent to which learning can, in an asymptotic sense, *undo* the effects of agents’ cognitive limitations.²¹ For example, a commonly-studied question in this literature is whether learning will in the long run lead to convergence to Nash equilibrium.

Perhaps the closest recent paper to ours is Barberis, Shleifer and Vishny (1998), hereafter BSV. Like we do, BSV consider agents who attempt to learn, but who are restricted to updating over a class of incorrect models. In their setting, the models are specifically about the persistence of the earnings process—one model is that shocks to earnings growth are relatively permanent, while another model is that these shocks are more temporary in nature.²² BSV’s conclusions about under- and overreaction to earnings news then follow directly from the mistakes that agents make in estimating persistence.

In our theory, the notion of a model is considerably more abstract: a model is *any construct* that implies that one sort of information is more useful for forecasting than another. Thus a model can be a metaphor like “Amazon is just another Barnes&Noble”, which might imply that it is particularly important to study Amazon’s gross margins. Or alternatively, a

²⁰Early contributions to the learning-in-games literature include Robinson (1951), Miyasawa (1961), and Shapley (1964). For a survey of more recent work, see Fudenberg and Levine (1998).

²¹ A similar comment can be made about the literature that asks whether learning by boundedly rational agents will lead to convergence to rational-expectations equilibria. See, e.g., Cyert and DeGroot (1974), Blume, Bray and Easley (1982), and Bray and Savin (1986).

²² In BSV, agents put zero weight on the model with the correct persistence parameter. One might argue that this assumption is hard to motivate, since the correct model is no more complicated or unnatural than the incorrect models that agents entertain. By contrast, in our setting, the correct multivariate model *is* more complicated than the simple univariate models that agents actually update over.

model can be “Company X seems a lot like Tyco”, which might suggest looking especially carefully at those footnotes in Company X’s annual report where relocation loans to executives are disclosed. We view it as a strength of our approach that we are able to obtain a wide range of empirical implications without having to spell out such details.

The representative-agent/model-selection version of our theory is also reminiscent of Mullainathan’s (2000) work on categorization. Indeed, our notion that individual agents practice model selection—instead of Bayesian model averaging—is essentially the same as Mullainathan’s rendition of categorization: “choosing a category which best fits the given data...instead of summing over all categories as the Bayesian would...” In spite of this apparent similarity, however, it is important to reiterate that our main empirical predictions do not come from a discrete category-switching mechanism as in Mullainathan (2000), but rather from the fact that agents restrict their updating to the class of simple models.

VI. Conclusions

This paper can be seen as an attempt to integrate learning considerations into a behavioral setting where agents are predisposed to using overly simplified forecasting models. The key assumption underlying our approach is that agents update only over the class of simple models, and place zero weight on the correct, more complicated model of the world. As we have demonstrated, this assumption yields a fairly rich set of empirical implications. Moreover, these implications seem to be robust to aggregation. That is, they come through either when there is a single representative agent who practices model selection, or when there is a market comprised of heterogeneous agents, in which case the market can be said to practice a form of model averaging.

While we have chosen to flesh out these implications in the specific context of the stock market, we do not at all mean to suggest that this is the only—or even the most—interesting application of our theory. Rather, we have focused on the case of the stock market because this seemed like a good way of turning the theory’s general content into a set of relatively precise empirical predictions that could be readily taken to the data. But to mention just one of many examples, it would seem that the basic ideas that we have developed could also be useful in thinking about, say, the ways in which employers go about the process of evaluating prospective job candidates. One task for future work is to map out some of these other applications of the theory in greater detail.

References

Abelson, R.P., 1978, *Scripts*, Invited address to the Midwestern Psychological Association, Chicago.

Asness, Clifford S., 1997, The interaction of value and momentum strategies, *Financial Analysts Journal* (March/April), 29-39.

Barberis, Nicholas, Andrei Shleifer and Robert W. Vishny 1998. A model of investor sentiment, *Journal of Financial Economics* 49, 307-343.

Becker, Holly, 2001, Amazon.com: The good, the bad and the ugly, *Lehman Brothers Equity Research*, p. 3 (March 2).

Blanchard, Olivier J. and Mark W. Watson, 1982, Bubbles, rational expectations, and financial markets, in Paul Wachtel (Ed.), *Crises in Economic and Financial Structure*, Lexington MA: Lexington Books, 295-315.

Blodgett, Henry, 1999, Still long-term upside..., *Merrill Lynch Equity Research*, p. 3 (March 9).

Blodgett, Henry, 2000, Soft Q2; lowering rating and reducing revenue and EPS estimates, *Merrill Lynch Equity Research*, p.1 (July 27).

Blume, Lawrence, Margaret M. Bray and David Easley, 1982, Introduction to stability of rational expectations equilibrium, *Journal of Economic Theory* 26, 313-317.

Bray, Margaret M. and N.E. Savin, 1986, Rational expectations equilibrium, learning and model specification, *Econometrica* 54, 1129-1160.

Bruner, J.S., 1957, Going beyond the information given, in H. Gulber and others (Eds.), *Contemporary Approaches to Cognition*, Cambridge, Mass: Harvard University Press.

Bruner J.S. and Postman, Leo, 1949, On the perception of incongruity: A paradigm, *Journal of Personality* 18, 206-223.

Chen, Joseph, Harrison Hong and Jeremy C. Stein, 2001, Forecasting crashes: Trading volume, past returns and conditional skewness in stock prices, *Journal of Financial Economics* 61, 345-381.

Cohen, Jonathan, 1998, Further thoughts on Amazon.com's operating model, *Merrill Lynch Equity Research*, p. 2 (September 15).

Combes, Genni, 2000, Commitment to narrow losses in 2000: Driving scale and operating efficiency, *Hambrecht and Quist Equity Research*, p. 4 (February 3).

Courtney, Tom, 1999, Thinking outside the big box superstore—A white paper on the internet retail revolution, *Banc of America Securities* (August).

Cyert, Richard M. and Morris H. DeGroot, 1974, Rational expectations and Bayesian analysis, *Journal of Political Economy* 82, 521-536.

Daniel, Kent D., David Hirshleifer and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and over-reactions, *Journal of Finance* 53, 1839-1885.

D'Eathe, Sara, 2001, Stagnant core growth and high fulfillment costs leave us cautious, *Thomas Weisel Partners Equity Research*, p. 2 (April 25).

Ehrens, Scott, 1999, ROIC is key (not the gross margin), *Bear Stearns Equity Research*, p. 3 (February 12).

Fama, Eugene F. and Kenneth R. French, 1992, The cross-section of expected returns, *Journal of Finance* 47, 427-65.

Farley, Sara, 2000, Still the leader, but risks outweigh upside potential, *Paine Webber Equity Research*, p. 2 (February 23).

Fiske, Susan T. and Shelley E. Taylor, 1991, *Social Cognition* (2nd ed.), New York: McGraw Hill, Inc.

Fudenberg, Drew and David K. Levine, 1998, Learning in games: Where do we stand?, *European Economic Review* 42, 631-639.

Harrison, J. Michael and David M. Kreps, 1978, Speculative investor behavior in a stock market with heterogeneous expectations, *Quarterly Journal of Economics* 93, 323-336.

Hirshleifer, David and Siew Hong Teoh, 2002, Limited attention, information disclosure and financial reporting, Ohio State University working paper.

Hong, Harrison and Jeremy C. Stein, 1999, A unified theory of underreaction, momentum trading and overreaction in asset markets, *Journal of Finance* 54, 2143-2184.

Hong, Harrison and Jeremy C. Stein, 2003, Differences of opinion, short-sales constraints and market crashes, *Review of Financial Studies* 16, 487-525.

Jegadeesh, Narasimhan and Titman, Sheridan, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 93-130.

Kahneman, Daniel and Amos Tversky, 1973, On the psychology of prediction, *Psychological Review* 80, 237-251.

Kandel, Eugene and Neil D. Pearson, 1995, Differential interpretation of public signals and trade in speculative markets, *Journal of Political Economy* 103, 831-72.

Kuhn, Thomas S., 1962, *The Structure of Scientific Revolutions*, Chicago: University of Chicago Press.

Kyle, Albert S. and F. Albert Wang, 1997, Speculation duopoly with agreement to disagree: Can overconfidence survive the market test? *Journal of Finance* 52, 2073-2090.

Lakonishok, Josef, Andrei Shleifer and Robert W. Vishny, 1994, Contrarian investment, extrapolation and risk, *Journal of Finance* 49, 1541-1578.

Lee, Charles and Bhaskaran Swaminathan, 2000, Price momentum and trading volume, *Journal of Finance* 55, 2017-2069.

Lord, C., L. Ross and M.R. Lepper, 1979, Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence, *Journal of Personality and Social Psychology* 37, 2098-2109.

Meeker, Mary, 1998, AOL—Part Deux! *Morgan Stanley Dean Witter Equity Research*, p.1-2 (August 4).

Miller, Edward, 1977, Risk, uncertainty, and divergence of opinion, *Journal of Finance* 32, 1151-1168.

Miyasawa, K., 1961, On the convergence of learning processes in a 2x2 non-zero sum game, Research Memorandum No. 33, Princeton University.

Morris, Stephen, 1996, Speculative investor behavior and learning, *Quarterly Journal of Economics* 111, 1111-1133.

Mullainathan, Sendhil, 2000, Thinking through categories, MIT working paper.

Nisbett, Richard and Lee Ross, 1980, *Human Inference: Strategies and Shortcomings of Social Judgment*, Englewood Cliffs, NJ: Prentice-Hall.

Odean, Terrance, 1998, Volume, volatility, price and profit when all traders are above average, *Journal of Finance* 53, 1887-1934.

Rabin, Matthew and Joel L. Schrag, 1999, First impressions matter: A model of confirmatory bias, *Quarterly Journal of Economics* 114, 37-82.

Robinson, Julia, 1951, An iterative method of solving a game, *Annals of Mathematics* 24, 296-301.

Schank, R. and R.P. Abelson, 1977, *Scripts, Plans, Goals and Understanding: An Introduction into Human Knowledge Structures*, Hillsdale, N.J.: Lawrence Erlbaum.

Scheinkman, Jose and Wei Xiong, 2003, Overconfidence and speculative bubbles, *Journal of Political Economy*, forthcoming.

Shapley, L.S., 1964, Some topics in two-person games, in M. Drescher, L.S. Shapley and A.W. Tucker (Eds.), *Advances in Game Theory*, Annals of Mathematics Study 52, Princeton NJ : Princeton University Press.

Simon, Herbert A., 1982, *Models of Bounded Rationality: Behavioral Economics and Business Organizations*, Vol. 2, Cambridge, MA: MIT Press.

Sims, Christopher A., 2003, Implications of rational inattention, *Journal of Monetary Economics* 50, 665-690.

Taylor, S.E. and J.C. Crocker, 1980, Schematic bases of social information processing, in E.T. Higgins, P. Herman and M.P. Zanna (Eds.), *The Ontario Symposium on Personality and Social Psychology* (Vol. 1), Hillsdale, N.J.: Lawrence Erlbaum.

Tversky, Amos and Daniel Kahneman, 1973, Availability: A heuristic for judging frequency and availability, *Cognitive Psychology* 5, 207-232.

Varian, Hal R., 1989, Differences of opinion in financial markets, in CC. Stone (Ed.), *Financial Risk: Theory, Evidence, and Implications: Proceedings of the 11th Annual Economic Policy Conference of the Federal Reserve Bank of St. Louis*, Boston MA: Kluwer Academic Publishers, pp. 3-37.

Table 1: Numerical Simulations

This table reports results from simulations under three cases: i) no learning; ii) learning with model selection; and iii) learning with model averaging. The number of time periods is $T=2,000$. The variances of the shocks are set to $v_a=v_b=v_\varepsilon=0.0001$. The autocorrelation coefficient of the processes A_t and B_t is set to $\rho=0.85$. And $r=0.03$, $h=0.051$, $\pi_A = \pi_B = 0.95$ and $\gamma=1$. We generate $N=1,000$ different time series of stock prices, and calculate a variety of statistics, based on averages across the simulations. ShiftProb is the number of paradigm shifts divided by T . Volatility is the square root of $E[(R_{t+1})^2]$. β^{MOM} is the coefficient in a regression of R_{t+1} on $R_{t-3,t}$ where $R_{t-3,t}$ is the cumulative return from $t-3$ to t inclusively. β^{VALUE} is the coefficient in a regression of R_{t+1} on P_t . Expected Return/Glamour is $E[R_{t+1} | P_t > 0]$. Expected Return/Glamour/Bad News is $E[R_{t+1} | P_t > 0, z_{t-3,t}^* < 0]$, where $z_{t-3,t}^*$ is the cumulation of the dividend-surprise components of returns from $t-3$ to t inclusively. Volatility/Glamour/Bad News is the square root of $E[(R_{t+1})^2 | P_t > 0, z_{t-3,t}^* < 0]$. Skewness/Glamour/Bad News is $E[(R_{t+1})^3 | P_t > 0, z_{t-3,t}^* < 0]$. ShiftProb/Glamour/Bad News is the probability of a shift in period $t+1$ given $P_t > 0$ and $z_{t-3,t}^* < 0$. Corr(A,B)/Glamour/Bad News is the correlation of A_t and B_t conditional on $P_t > 0$ and $z_{t-3,t}^* < 0$. Expected Return/Glamour/Bad Returns is $E[R_{t+1} | P_t > 0, R_{t-3,t} < 0]$. Volatility/Glamour/Bad Returns is the square root of $E[(R_{t+1})^2 | P_t > 0, R_{t-3,t} < 0]$. Skewness/Glamour/Bad Returns is $E[(R_{t+1})^3 | P_t > 0, R_{t-3,t} < 0]$. ShiftProb/Glamour/Bad Returns is the probability of a shift in period $t+1$ given $P_t > 0$ and $R_{t-3,t} < 0$. Corr(A,B)/Glamour/Bad Returns is the correlation of A_t and B_t conditional on $P_t > 0$ and $R_{t-3,t} < 0$. In Panels B and C, we change ρ to 0.90 and 0.95 respectively.

Panel A: $\rho=0.85$

	1. No Learning	2. Learning: model selection	3. Learning: model averaging
ShiftProb	NA	0.0310	NA
Volatility	0.0596	0.0734	0.0784
β^{MOM}	0.0547	0.0315	0.0074
β^{VALUE}	0	-0.0286	-0.0357
Expected Return/Glamour	0	-0.0034	-0.0036
Expected Return/Glamour/Bad News	-0.0109	-0.0150	-0.0174
Volatility/Glamour/Bad News	0.0596	0.0818	0.0915
Skewness/Glamour/Bad News	0	-0.0009	-0.0012
ShiftProb/Glamour/Bad News	NA	0.0484	NA
Corr(A,B)/Glamour/Bad News	NA	-0.4250	-0.5553
Expected Return/Glamour/Bad Returns	-0.0114	-0.0145	-0.0066
Volatility/Glamour/Bad Returns	0.0596	0.0799	0.0831
Skewness/Glamour/Bad Returns	0	-0.0008	-0.0006
ShiftProb/Glamour/Bad Returns	NA	0.0476	NA
Corr(A,B)/Glamour/Bad Returns	NA	-0.3173	-0.4571
<i>Note: $k=5.56$, Rational-Expectations Volatility=0.0792</i>			

Panel B: $\rho=0.90$

	1. No Learning	2. Learning: model selection	3. Learning: model averaging
ShiftProb	NA	0.0392	NA
Volatility	0.0809	0.1094	0.1204
β^{MOM}	0.0505	-0.0087	-0.0415
β^{VALUE}	0	-0.0336	-0.0593
Expected Return/Glamour	0	-0.0071	-0.0121
Expected Return/Glamour/Bad News	-0.0148	-0.0277	-0.0402
Volatility/Glamour/Bad News	0.0809	0.1240	0.1492
Skewness/Glamour/Bad News	0	-0.0054	-0.0063
ShiftProb/Glamour/Bad News	NA	0.0560	NA
Corr(A,B)/Glamour/Bad News	NA	-0.4633	-0.5965
Expected Return/Glamour/Bad Returns	-0.0128	-0.0197	-0.0198
Volatility/Glamour/Bad Returns	0.0809	0.1204	0.1342
Skewness/Glamour/Bad Returns	0	-0.0037	-0.0036
ShiftProb/Glamour/Bad Returns	NA	0.0550	NA
Corr(A,B)/Glamour/Bad Returns	NA	-0.2577	-0.4148
<i>Note: $k=7.69$, Rational-Expectations Volatility=0.1092</i>			

Panel C: $\rho=0.95$

	1. No Learning	2. Learning: model selection	3. Learning: model averaging
ShiftProb	NA	0.0469	NA
Volatility	0.1295	0.2066	0.2345
β^{MOM}	0.0442	-0.0509	-0.0903
β^{VALUE}	0	-0.0361	-0.0658
Expected Return/Glamour	0	-0.0198	-0.0380
Expected Return/Glamour/Bad News	-0.0231	-0.0639	-0.1010
Volatility/Glamour/Bad News	0.1295	0.2538	0.3051
Skewness/Glamour/Bad News	0	-0.0552	-0.0653
ShiftProb/Glamour/Bad News	NA	0.0658	NA
Corr(A,B)/Glamour/Bad News	NA	-0.5818	-0.6238
Expected Return/Glamour/Bad Returns	-0.0148	-0.0357	-0.0596
Volatility/Glamour/Bad Returns	0.1295	0.2326	0.2749
Skewness/Glamour/Bad Returns	0	-0.0374	-0.0475
ShiftProb/Glamour/Bad Returns	NA	0.0650	NA
Corr(A,B)/Glamour/Bad Returns	NA	-0.2727	-0.3484
<i>Note: k=12.50, Rational-Expectations Volatility=0.1771</i>			

Figure 1: Stock Price History for Amazon.com, May 1997-January 2003

