

FINANCE RESEARCH SEMINAR SUPPORTED BY UNIGESTION

“A Dynamic Model of Characteristic- Based Return Predictability”

Prof. Aydogan ALTI

University of Texas, McCombs School of Business

Abstract

We present a model where the “climate of disruptive innovation,” which determines the arrival rate of new projects and the exit rate of existing businesses, is a source of systematic risk that influences the returns of portfolios sorted on value, profitability, and asset growth. These characteristic-sorted portfolios generate abnormal returns along finite sample paths too frequently, relative to the rational benchmark, when investors exhibit biases in processing information about the disruption climate. In simulations of the calibrated model we quantify the ex-ante likelihood of observing the historical evidence on characteristic-based return predictability.

Friday, March 24, 2017, 10:30-12:00

Room 126, Extranef building at the University of Lausanne

A Dynamic Model of Characteristic-Based Return Predictability

AYDOĞAN ALTI and SHERIDAN TITMAN*

Current Version: March 2017

ABSTRACT

We present a model where the “climate of disruptive innovation,” which determines the arrival rate of new projects and the exit rate of existing businesses, is a source of systematic risk that influences the returns of portfolios sorted on value, profitability, and asset growth. These characteristic-sorted portfolios generate abnormal returns along finite sample paths too frequently, relative to the rational benchmark, when investors exhibit biases in processing information about the disruption climate. In simulations of the calibrated model we quantify the ex-ante likelihood of observing the historical evidence on characteristic-based return predictability.

* University of Texas at Austin. An earlier draft of the paper was circulated under the title “Creative Destruction, Investor Beliefs and the Evolution of Stock Returns.” We benefited from comments and suggestion by Zhanhui Chen, Travis Johnson, Mao Lei, David Solomon, Kristian Rydqvist, and participants in presentations of the paper at the 2016 Rising Stars Conference at Fordham University, Workshop in Memory of Rick Green at Carnegie Mellon University, 2016 Asian Bureau of Finance and Economic Research Conference in Singapore, Asset Pricing Conference in Honor of John Wei at HKUST, 2016 UBC Summer Finance Conference, Emory University, SAIF, University of Colorado Boulder and Denver, University of Georgia, University of Texas at Austin, and Vienna University of Economics and Business.

Since the late 1970s,¹ financial economists have identified a variety of firm characteristics and financial variables that seem to explain the cross-sectional pattern of stock returns.² Although many of the anomalies relate to short-term return patterns (e.g., momentum and return reversals) many relate to fundamental firm characteristics like valuation ratios, profitability and asset growth rates, which are the focus of this paper. Historically, market-neutral portfolios that are designed to exploit these fundamentals-based anomalies have exhibited extremely high Sharpe ratios – at least as high as the Sharpe ratio of the market.

These historical returns patterns present a considerable challenge. Rational asset pricing models, which require strong assumptions about preferences to explain the market risk premium, are not likely to explain even higher Sharpe ratios on market-neutral portfolios.³ Behavioral models have more degrees of freedom, but existing behavioral models do not deliver a clear link between firm characteristics and returns and are not designed to assess quantitative magnitudes.

This paper develops and calibrates a dynamic model that links characteristic-based return predictability to the mispricing of a systematic risk factor that affects the cross section of firms' fundamentals. In contrast to the existing literature we do not ask whether *expected* Sharpe ratios match historical Sharpe ratios; by design, the *unconditional* expected return associated with the systematic risk factor is zero in our model. However, since the mispricing in the model is slow to correct, excess returns are serially correlated, which causes long-term returns to be too volatile relative to short-term returns. This, in turn, increases the dispersion of the distribution of realized Sharpe ratios, increasing the likelihood of extreme realizations. In model simulations we quantify

¹ See the 1978 special issue of the *Journal of Financial Economics* (Vol. 6, Issues 2-3) on anomalous evidence regarding market efficiency.

² Indeed, in a recent paper, McLean and Pontiff (2016) identify 97 variables that have been shown to predict cross-sectional stock returns.

³ MacKinlay (1995) was the first published paper that argued that the Sharpe ratios of characteristic-sorted portfolios are simply too high to be consistent with ex-ante rational expectations.

this effect and generate Sharpe ratios of characteristic-sorted portfolios that match or exceed historical Sharpe ratios in up to 10% of the sample paths, which contrasts sharply with the extremely low likelihoods obtained under full rationality. Our analysis thus suggests that we should be less confident about the unconditional expected returns of these characteristic-sorted portfolios being positive, and similarly less confident about the high Sharpe ratios observed over the past 50 years repeating over the next 50 years.

Firms in our model, which can be viewed as combinations of assets in place and growth opportunities, differ in their access to new growth opportunities – growth firms are endowed with new projects every period while value firms simply harvest the profits from their existing opportunities – and their histories – whether or not they were endowed with projects in the past. The emergence of these growth opportunities, as well as the demise of existing businesses, are determined by a systematic factor that we refer to as the climate of disruptive innovation. A favorable disruption climate increases the arrival rate of new projects, which benefits the young growth firms, but because these new projects compete with existing businesses, a favorable disruption climate harms the profits of assets in place, and is thus detrimental to the value firms. The model can thus be described as a Schumpeterian model of creative destruction, where innovation creates losers as well as winners.

Investors receive noisy indicators of the disruption climate, and thus form expectations that in general contain estimation errors. Because of these errors, even fully rational investors are sometimes too optimistic and sometimes too pessimistic about the rate of future disruptive innovations. If, however, investors correctly assess the precision of their signals and process new information rationally, then sample paths that reject market efficiency at the 10% level occur about 10% of the time. In other words, parameter uncertainty does not by itself lead to biased

inferences – some degree of irrationality is needed to explain the observed asset pricing anomalies.

Investor perceptions in our model may be biased for two reasons that have been previously discussed in the literature. First, since investors learn about the disruption climate from soft information, like news reports and expert opinions, they may find it difficult to assess the precision of their signals. We assume that if investors are biased in their assessment of these precisions, they will tend to believe the signals are more precise than they really are. In other words, investors will tend to be overconfident about their abilities to evaluate soft information. This bias, which is the main focus of our analysis, generates return persistence and increases the dispersion of the Sharpe ratio distribution. We also consider a second source of bias that comes from the fact that the initial beliefs of investors may be wrong. This bias, which can affect the mean of the return distribution, is motivated by Shiller (2000) and others who suggest that investors are overly optimistic about the commercialization potential of technical innovations.

Our model, which is calibrated under alternative assumptions about investor perceptions, is able to generate many of the qualitative features of the data. In particular, along sample paths where investors tend to be negatively surprised about the disruption rate, value, profitable, and low asset growth firms beat growth, unprofitable, and high asset growth firms, respectively. The model's comparative statics are generally consistent with our intuition about the economic structure that can generate the observed returns of characteristic-sorted portfolios. In particular,

- Increasing overconfidence about the precision of the soft signal increases the dispersion of the distribution of Sharpe ratios that are generated in the model.

- The persistence of the disruption climate is also important. Biased perceptions about disruption cannot persist if the climate materially changes from month to month, so one needs a reasonably stable environment for mistakes about the disruptive potential of new technology to create significant persistence in returns.
- Greater optimism shifts the distribution of Sharpe ratios. Our simulations, however, demonstrate that if investors have an unbiased assessment of the soft signal's precision, they learn rationally and the effect of being initially overly optimistic dissipates quickly.

Finally, we examine how quantitative investment strategies that are designed to detect and exploit mispricing perform in this setting. This exercise is of interest for two reasons. First, to our knowledge, we are providing the first analysis of the returns that can be generated by what we are describing as “quant” investors. The second is that this exercise provides a sanity check on our model's calibration. One might reject our contention that we are considering a setting where investors make plausible mistakes, if the setting allows quant investors to generate very high Sharpe ratios on average.

Our thought experiment considers different types of quant investors. At one extreme we consider a fully rational investor with complete knowledge of the structure of our economy. Our simulations imply that such an investor could be expected to obtain a yearly Sharpe ratio of 0.41, which is very high for a market-neutral portfolio. At the other extreme we consider a quant investor that is informed only by the past returns of characteristic-sorted portfolios. Such an investor is expected to generate a more modest yearly Sharpe ratio of about 0.22 when utilizing recent data samples and even less with longer data samples.

Although our overall approach is substantially different, we build on a number of contributions to the asset pricing literature. The analysis in this paper is particularly related to asset pricing models with parameter uncertainty. Most notably, Lewellen and Shanken (2002) show, within a setting where rational Bayesian investors learn about expected cash flows, that returns will appear to be predictable conditional on the realization of the expected cash flow, even when they are not at all predictable ex ante. Similarly, returns in our model appear to be predictable along some sample paths. Part of the contribution of our model is that it allows us to assess the likelihood of observing such sample paths. Although their focus is on the market return, Pastor and Stambaugh (2012) also analyze long-term return variances in a model with parameter uncertainty and changing expected rates of return. Similar to our paper, they make the point that changes in expected returns can increase the volatility of long-term return realizations. However, in contrast to our model, which analyzes the pricing of uncertain cash flow streams, Pastor and Stambaugh consider exogenously-specified return processes.⁴

Our analysis is also related to the behavioral finance literature, and in particular to Daniel, Hirshleifer, and Subrahmanyam (DHS) (1998, 2001), which describe a link between the value effect and the tendency of investors to be overconfident about the precision of their private information.⁵ We also focus on overconfidence, but our channel generating mispricing is very different. In the DHS papers, firms are essentially identical, and the value effect is mechanically generated from the fact that overpriced stocks tend to have high prices relative to fundamentals.

⁴ Other related papers include Timmermann (1993, 1996) and Pastor and Veronesi (2003, 2006). Timmermann (1993, 1996) analyzes models where investors use Bayes' rule to estimate unknown parameters but value assets without taking into account estimation error, resulting in predictable returns and excess volatility. Pastor and Veronesi (2003, 2006) also analyze learning effects on asset prices, though their focus is not on return predictability. The Pastor-Veronesi models illustrate how uncertainty about future growth rates may rationalize high and volatile valuations, especially for your businesses such as the technology firms of the late 1990s.

⁵ Other important papers in this literature include Barberis, Shleifer and Vishny (1998), which considers behavioral biases that influence how investors estimate the persistence of earnings shocks, and Hong and Stein (1999), which considers the effects of positive-feedback traders and traders who ignore the information embedded in market prices.

In contrast, the firms in our model differ in fundamental ways, i.e., the growth firms have ongoing new investment opportunities and the value firms do not, and it is these fundamental differences that lead to cross-sectional differences in their exposures to sources of systematic risk. As a result, our model addresses differences in rates of return associated with characteristics like profitability and asset growth rates as well as price scaled characteristics like the market-to-book ratio. In addition, we develop a dynamic model that can be calibrated and used to simulate returns that can be compared to the actual time-series pattern of stock returns.

There is also a behavioral literature that explores how fluctuations in investor sentiment can induce covariation amongst stocks with common characteristics.⁶ In our model, overconfident investors tend to over-react to soft information about disruption shocks, and in doing so they induce excess (relative to the fully rational case) covariation amongst stocks with similar characteristics. In this sense, our model endogenously generates what looks like a sentiment factor.

This research complements recent work by Gârleanu, Kogan, and Panageas (2012) and Kogan, Papanikolaou, and Stoffman (2015), which examine how the innovation process can generate sources of systematic risk that affect the prospects of different firms differently. In these models growth firms earn low expected returns as they constitute a hedge against the displacement risk brought about by technological progress. These models, which assume full rationality, require fairly strong risk preferences to rationalize the historically observed return patterns. Although we solve our model with risk neutral preferences, effectively shutting down

⁶ See Baker and Wurgler (2006) for an empirical analysis of the impact of time-varying investor sentiment on characteristic-sorted portfolio returns and Kozak, Nagel, and Santosh (2015) for a more complete discussion of how investor sentiment can influence systematic risk factors.

the risk aversion channel, we can envision a model that accounts for risk preferences as well as slow learning that better explains the historical return patterns.⁷

It should be stressed that both the behavioral and the rational approaches in the existing literature address the historical characteristic-sorted portfolio returns with models that focus on *expected* Sharpe ratios. In contrast, our focus is on behavioral biases that can increase the *dispersion* of the distribution of possible Sharpe ratios. In other words, we attempt to rationalize historical Sharpe ratios as tail outcomes that are much more likely to occur when perceptions are biased, rather than as expected outcomes.

The remainder of the paper is organized as follows. Section I presents the model. Section II describes the historical data sample, the model calibrations and simulations procedures. Section III presents the results from calibrated model analyses. Section IV concludes the paper. Appendix A contains the technical derivations.

I. THE MODEL

As we mentioned in the introduction, the historical evidence provides a significant challenge. The returns of various characteristic-sorted portfolios appear to be both too large and too persistent. In this section we present a model that can be used to gauge quantitative as well as qualitative relationships. The model, which is designed to generate cross-sectional differences in firm characteristics, such as valuation ratios, asset growth rates, and profitability, is then applied to explore the relationship between these characteristics and returns.

⁷ A number of earlier papers in the literature provide risk-based explanations for the value premium. Examples are Berk, Green, and Naik (1999), Carlson, Fisher, and Giammarino (2004), and Zhang (2005). These papers, by specifying exogenous pricing kernels that are calibrated with a high price of risk, also effectively assume extreme risk aversion.

A. Model setup

Time is continuous and denoted by t . Investors are risk neutral, and they discount cash flows at a constant rate r . The investors also have the same beliefs about model parameters and observe the same information, which implies that asset prices are effectively set by a price-taking representative agent. The economy is populated by a continuum of firms and we use the superscript i to denote a generic individual firm. Financing is frictionless and the Modigliani-Miller Theorem holds so we can assume all firms are equity financed without loss of generality.

Firms derive their values from the projects that they initiate, which are infinitesimal investment opportunities that arrive continuously over time. Once a project is initiated, it generates cash flows until it becomes obsolete and is terminated. Specifically, a new project requires a capital investment k_z and generates deterministic cash flows at rate $a_z \times k_z$ until it is terminated, where a_z represents the project's return on investment.⁸ When a project is terminated, a fraction α of the initial investment k_z is recovered.⁹ The arrival and the termination rates of projects are determined by an economy-wide state variable that we describe below. The net cash flows of the firm, which include the cash flows from the projects, the costs of initiating new projects, and the proceeds from liquidations, are immediately paid out to shareholders.¹⁰

Firms in this model differ for two reasons. The first is that they have different histories,

⁸ The capital investment k_z and the return on investment a_z may depend on the firm-specific growth state z ; see below. Our assumption of projects with deterministic cash flows allows us to focus on the arrival and the termination of projects as the primary sources of risk. A more general version of the model could feature project cash flows that are subject to additional risk factors.

⁹ For simplicity we assume that capital does not depreciate. An alternative interpretation is that capital depreciates but has to be replenished for the project to be operational; in this interpretation, the cost of depreciation is implicit in the project return a_z .

¹⁰ By assuming an exogenous process governing project arrivals and terminations, we abstract from the possibility that firms' real investment choices are influenced by investor beliefs. This assumption allows us to focus on the pricing of a given set of assets. Yet the feedback from asset prices to investment choices may be relevant for some return anomalies, especially those that relate to firm fundamentals. See Alti and Tetlock (2014) for a quantitative analysis of how firms' investment decisions may amplify the impact of biased investor beliefs on return predictability.

i.e., they initiated different projects in the past. The second is that they inhabit different states, which determine the new projects they receive. Specifically, a firm is in one of three states at any given time: the *early growth* state EG , the *mature growth* state MG , and the *no growth* state NG . Let $z_t^i \in \{EG, MG, NG\}$ describe the state of firm i at time t . *New firms* are born with identical initial conditions into the early growth state and they can transition into the mature growth state and the no growth state over time.

In our calibrated model, the firm has access to new projects in the early and the mature growth states, but not in the no growth state. Specifically, we set $k_{EG} = k_{MG} = k$ and $a_{EG} > a_{MG}$. Thus, the projects that arrive in the mature growth state require the same initial capital investment k as in the early growth state, but are less profitable. The assumption that firms in the no growth state receive no projects can be stated as $k_{NG} = a_{NG} = 0$.

Let f_t^i denote the firm's *profitability*, defined as the rate at which the cash flow from the firm's active projects are generated, and K_t^i denote the firm's *capital stock*, which is the total capital investment incurred for the active projects. A new firm i , which is born at time t into the early growth state, has an initial capital stock that is normalized to $K_t^i = 1$ and a profitability of $f_t^i = 0$.¹¹

After being born into the early growth state, the firm's state z_t^i evolves as a continuous-time Markov process with sequential jumps. Specifically, the firm transitions from the early growth state to the mature growth state with Poisson intensity q_{EG} , and from the mature growth

¹¹ The assumption that firms are born with an unproductive unit of capital is motivated by the presence of firms with valuable growth opportunities but little or no profits in the data. The specific value chosen here, zero initial profitability, does not affect our results in any material way; what is important is that the model includes growth firms with low profits.

state to the no growth state with Poisson intensity q_{MG} . A firm in the no growth state dies and leaves the firm population with Poisson intensity q_{NG} . Each firm that dies is replaced by a new firm that is born into the early growth state, as described above. When a firm dies, its owners receive the market value of its active projects as a liquidating dividend.¹² The transition rates described above imply that firms spend $1/q_{EG}$ years on average in the early growth state, $1/q_{MG}$ years on average in the mature growth state, and $1/q_{NG}$ years on average in the no growth state. Thus, the average life expectancy of a firm is $1/q_{EG} + 1/q_{MG} + 1/q_{NG}$ years.

As long as the firm is alive, its capital stock K_t^i and its profitability f_t^i evolve according to the following laws of motion:

$$dK_t^i = k_z (dt + dM_t) - \lambda K_t^i (dt + dM_t), \quad (1)$$

$$df_t^i = a_z k_z (dt + dM_t) - \lambda f_t^i (dt + dM_t). \quad (2)$$

In Equations (1) and (2), the z subscript refers to the firm-specific growth state. The term dM_t represents a systematic *disruption rate*, which is a persistent process with a long-term mean of zero. Before we describe the specific stochastic process governing dM_t , we will first explain how it interacts with the firm-specific state variables in Equations (1) and (2).

The first terms in Equations (1) and (2) capture the arrival of new projects. Recall that firms in the early and the mature growth states receive projects that each require a capital

¹² Thus, a firm's death in our model resembles an asset sale to an entity outside the publicly traded corporate sector.

investment of $k_z = k$ and add $a_z \times k_z$ to the firm's profitability. The rate at which new projects arrive is stochastic and is represented by the term $dt + dM_t$. Thus, over an instantaneous time period dt the firm receives dt projects on average (i.e., one project per unit of time), with more or less new projects arriving depending on the realization of the disruption rate dM_t .

The second terms in Equations (1) and (2) reflect the termination of active projects. Over a given period of time, a fraction of the firm's active projects become obsolete and are liquidated. When a project is terminated the capital of the firm declines and the profitability of the firm declines by a proportional amount. Projects are terminated at an average rate $\lambda > 0$, which we assume is the same for all firms. As with the arrival of new projects, the realized termination rate depends on the disruption rate and is given by $\lambda(dt + dM_t)$.

Equations (1) and (2) illustrate how the exogenous growth state of the firm, i.e., early growth, mature growth or no growth, generates the endogenous state variables, capital stock K_t^i and profitability f_t^i . Transitions from one state to another, along with the termination rate of active projects, generate cross-sectional and time-series variations in firm size, profitability, and valuation ratios. Firms that have profitable active projects but are in the no growth state expect their size and profits to decline over time. Firms that have low current profitability but are in growth states expect the opposite. Thus, the model captures in a reduced-form way the Schumpeterian notion of creative destruction: profits are redistributed from established firms to newcomers and from old to new technologies.

The economy-wide disruption rate dM_t , which determines the speed of this Schumpeterian reallocation process, is the main focus of the model. When dM_t is high, new projects are created faster and existing projects are destroyed faster. As a result, early-growth

firms benefit when dM_t is high and no-growth firms are hurt. Depending on parameters, mature-growth firms can either be helped or hurt by a more innovative environment that creates more disruption, since it can hurt their existing businesses while at the same time facilitating new projects.

The disruption rate, which is observable, is an exogenously specified process that has both persistent and transitory components. Specifically,

$$dM_t = \mu_t dt + \sigma_M d\omega_t^M, \quad (3)$$

where μ_t , what we have referred to as the *disruption climate*, is the persistent component of the disruption rate, and the Brownian process $d\omega_t^M$ is the transitory component. We envision the disruption climate μ_t as a slow-moving variable with a long-term mean that is normalized to be zero. Specifically, μ_t evolves according to

$$d\mu_t = -\rho_\mu \mu_t dt + \sigma_\mu d\omega_t^\mu. \quad (4)$$

Although investors observe the realized disruption rate dM_t , they cannot separately observe the persistent and the transitory components.¹³ They do, however, observe a *soft information signal* ds_t that reflects the state of technological progress, changes in the regulatory

¹³ Investors observe the realized disruption rate dM_t because each firm's changes in capital stock and profitability are observable and can be used to back out dM_t (see Equations 1 and 2). In a more general version of the model where change in profitability contains additional noise terms, a single firm's change in profitability would not perfectly reveal dM_t , but with a large cross section of firms investors would still be able to estimate it highly precisely.

environment, and other information that may help them predict the future evolution of the disruption climate. Specifically, investors observe

$$ds_t = \eta d\omega_t^\mu + \sqrt{1-\eta^2} d\omega_t^s, \quad (5)$$

where the parameter $\eta \in [0,1]$ is the signal's precision and the Brownian term $d\omega_t^s$ is the signal's noise. Higher values of η describe a more informative signal and thus less residual uncertainty about μ_t . In model calibrations, we consider the possibility that investors have biased perceptions about the precision of the soft information signal. Specifically, we analyze cases where overconfident investors believe the signal precision parameter to be $\eta_B > \eta$.

Investors in this model use the historical realizations of the disruption rate along with their soft information to learn about expected future disruption rates. We model the disruption rate dM_t in a way that reflects the learning features that we would like to analyze. For learning to be relevant, the disruption rate needs to have a persistent component that is not directly observable along with a transitory component. In other words, as expressed in Equation (3), the observed disruption rate equals the persistent component plus a transitory component that obscures the investor's inference problem. In principle, the persistent component μ_t could be an unknown constant μ ; however, when this is the case, learning effects vanish in the long run since investors eventually learn μ arbitrarily precisely. In our model, the persistent component μ_t changes over time. Investors learn about the current value of μ_t , but unobservable shocks to μ_t create an additional source of uncertainty. In the steady state, these two effects cancel out, and the estimation error investors face about μ_t remains constant over time.

The soft signal ds_t plays an important role in the model. The signal summarizes all the non-financial data that investors use to evaluate the current disruption climate. Investors' possibly biased perception of the signal's precision is the driver of the return predictability patterns that the model generates. The signal is assumed to be informative about the shocks to μ_t rather than the level of μ_t . This specification, which we take from Scheinkman and Xiong (2003), has two advantages. First, the signal has constant variance $\eta^2 + (\sqrt{1-\eta^2})^2 = 1$ regardless of the value of η . Thus, the specification permits biased investor beliefs about signal precision (i.e., $\eta_B \neq \eta$) that cannot be detected directly from the time-series variance of signal realizations. Second, the specification clearly delineates the two sources of information investors use to update their estimates. The signal is informative about shocks to μ_t whereas the realized disruption rate dM_t is informative about the current level of μ_t . Being orthogonal to each other, these two sources of information generate an economically meaningful two-factor structure for asset returns.

B. Information Processing

As discussed above, investors update their beliefs about the disruption climate μ_t based on two pieces of information, the realized disruption rate dM_t and the signal ds_t . In this subsection we characterize the steady state of the model in which the precision of the conditional estimate of μ_t is constant over time.

Let $\hat{\mu}_t$ denote investors' *expected disruption rate*, defined as the conditional estimate of μ_t given all available information at time t . Let γ denote the steady-state variance of the

estimation error $\hat{\mu}_t - \mu_t$. The law of motion of $\hat{\mu}_t$ is given by

$$d\hat{\mu}_t = -\rho_\mu \hat{\mu}_t dt + \sigma_\mu \eta ds_t + \frac{\gamma}{\sigma_M} d\bar{\omega}_t, \quad (6)$$

where

$$d\bar{\omega}_t \equiv \frac{dM_t - \hat{\mu}_t dt}{\sigma_M} \quad (7)$$

is a standard Brownian motion that corresponds to the surprise component of the realized disruption rate. Recall that the signal ds_t is also a standard Brownian motion by construction.

Therefore, $d\bar{\omega}_t$ and ds_t constitute the two sources of systematic risk that are orthogonal to each other.

The steady-state variance of the estimation error γ solves

$$\frac{\sigma_\mu^2}{2\rho_\mu} = \frac{1}{2\rho_\mu} \left(\sigma_\mu^2 \eta^2 + \frac{\gamma^2}{\sigma_M^2} \right) + \gamma. \quad (8)$$

The solution is given by

$$\gamma = \sigma_M \left(-\rho_\mu \sigma_M + \sqrt{\rho_\mu^2 \sigma_M^2 + (1 - \eta^2) \sigma_\mu^2} \right). \quad (9)$$

When investors have biased signal precision, the parameter η in Equations (6), (8), and

(9) is replaced by its biased counterpart $\eta_B > \eta$, which results in $\gamma_B < \gamma$. When this is the case, investors overestimate the precision of their soft information, which implies that they believe that their disruption rate estimate $\hat{\mu}_t$ is more precise than it actually is. Inspecting Equation (6), we see that this bias leads investors to place too much weight on their soft signal and too little weight on the realized disruption rate in updating $\hat{\mu}_t$.

C. Firm Valuation and Returns

We now turn to firms' valuations and their stock returns. To keep our discussion focused on the economic intuition we will present the relevant equations and provide their derivations in Appendix A.

The values of the firms can be expressed as the discounted value of their expected cash flows conditional on all available information:

$$V_t^i = E_t \left(\int_{u=t}^{\infty} e^{-r(u-t)} \left[f_u^i du + (\alpha \lambda K_u^i - k_z)(du + dM_u) \right] \right). \quad (10)$$

The cash flows in Equation (10) consist of three components, which include f_u^i , the profits accruing to firm i from its active projects, $\alpha \lambda K_u^i$, the capital recovered from terminated projects, and k_z , the outflows arising from capital investments for new projects given the firm-specific state $z = z_u^i$. Note that the profits from active projects accrue at an instantaneously deterministic rate, whereas capital flows for new and terminated projects are stochastic and determined by the economy-wide disruption rate dM_u .

The firm value in Equation (10) can be decomposed as follows:

$$V(z_t^i, K_t^i, f_t^i, \hat{\mu}_t) = f_t^i V_f(\hat{\mu}_t) + K_t^i V_K(\hat{\mu}_t) + V_{g,z}(\hat{\mu}_t), \quad (11)$$

where the functions $V_f(\hat{\mu}_t)$, $V_K(\hat{\mu}_t)$, and $V_{g,z}(\hat{\mu}_t)$ are solutions to a set of ordinary differential equations. The first two terms in Equation (11) respectively represent the present value of the cash flows from the firm's active projects and the value derived from the expected partial recovery of capital tied to those projects. Note that these two terms are linear in the firm's profitability f_t^i and its capital stock K_t^i . The third term, which is a function of the firm's idiosyncratic growth state z , reflects the NPV of future projects.¹⁴ When investors expect a relatively high degree of disruption (i.e., $\hat{\mu}_t$ is high), firms in the early and the mature growth states enjoy higher valuations of their growth opportunities (i.e., $V_{g,EG}(\hat{\mu}_t)$ and $V_{g,MG}(\hat{\mu}_t)$ are relatively high). In these same high $\hat{\mu}_t$ states, active projects are expected to be terminated sooner, which implies that the present values of firm cash flows $V_f(\hat{\mu}_t)$ are low and the values derived from the partial recovery of capital $V_K(\hat{\mu}_t)$ are high.

The firm's excess return (i.e., its realized rate of return in excess of the discount rate r) consists of three components, as expressed below:

¹⁴ The functions $V_{g,z}$ account for future transitions of the growth state and future projects' initial capital investments, profits, and eventual capital recoveries. Also, note that growth opportunities are worth zero in the no growth state, i.e., $V_{g,NG} \equiv 0$.

$$\begin{aligned}
dr_t^i - rdt = & \left[\frac{f_t^i V_f'(\hat{\mu}_t) + K_t^i V_K'(\hat{\mu}_t) + V_{g,z}'(\hat{\mu}_t)}{V_t^i} \right] \left(\sigma_\mu \eta ds_t + \frac{\gamma}{\sigma_M} d\bar{\omega}_t \right) \\
& + \left[\frac{(\alpha \lambda K_t^i - k_z) + (k_z - \lambda K_t^i) V_K(\hat{\mu}_t) + (a_z k_z - \lambda f_t^i) V_f(\hat{\mu}_t)}{V_t^i} \right] \sigma_M d\bar{\omega}_t + d\varepsilon_t^i.
\end{aligned} \tag{12}$$

The first two terms in Equation (12) characterize the exposure of the firm's return to the two systematic risk factors in the model, the disruption surprise $d\bar{\omega}_t$ and the signal ds_t . The last term $d\varepsilon_t^i$ is the firm's idiosyncratic return that is driven by growth state transitions.

A comparison of the first two terms in Equation (12) provide some intuition for the factor structure of asset returns in the model. The first term in Equation (12) reflects the Bayesian updating of the expected disruption climate $\hat{\mu}_t$, which is relevant for predicting subsequent disruption rates and thus valuing future cash flows. The second term in Equation (12) captures the *immediate impact* of the disruption process on the arrival and termination of projects. Note that both the disruption surprise $d\bar{\omega}_t$ and the signal ds_t contribute to the updating of $\hat{\mu}_t$, but only the disruption surprise $d\bar{\omega}_t$ has an immediate effect on the firm's projects. This asymmetry is what generates a two-factor structure in asset returns. Having experienced different histories and being in different growth states, firms differ in their relative sensitivities to the expected disruption climate versus the realized disruption rate. Thus, firms' relative exposures to the two systematic risk factors differ in the cross section.

Up to this point we have characterized firm values and returns conditional on the beliefs of investors, which may or may not be biased. In the rest of this section we consider the case where investors have biased beliefs, but characterize expected returns from the perspective of a

fully rational observer. In particular, we will be examining the link between the two systematic risk factors in the model, the signal ds_t and the disruption surprise $d\bar{\omega}_t$, and return predictability.

First, consider the signal ds_t . From Equation (12) we see that the impact of the signal ds_t on returns is proportional to η . This implies that if investors have biased beliefs about the signal's precision, i.e., $\eta_B > \eta$, the sensitivity of returns to the signal is amplified. This amplification effect, however, does not influence the conditional expected rates of return. This is because the signal ds_t has an expected value of zero regardless of either the perceived or the actual precision. Even with biased perceptions of the precision of the signal, investors are observing something that is a random walk so their expectation of what the signal will be in the next instance is unbiased.

In contrast to the signal, the disruption surprise $d\bar{\omega}_t$ can generate conditional return predictability when investors have biased beliefs. The reason is that historical mistakes that investors make in processing information generate a conditional drift in $d\bar{\omega}_t$ from the perspective of a rational observer. To illustrate formally, let $\hat{\mu}_t^B$ denote investors' estimate of μ_t in the case where they have a biased perception of the signal precision $\eta_B > \eta$. Let $\hat{\mu}_t^R$ denote the unbiased estimate a fully rational observer (i.e., someone who knows the true signal precision) would have given the same history. Similarly, let $d\bar{\omega}_t^B$ and $d\bar{\omega}_t^R$ denote the disruption surprise from the perspectives of biased investors and the rational observer, respectively. Substituting the new notation into Equation (7) yields

$$d\bar{\omega}_t^B \equiv \frac{dM_t - \hat{\mu}_t^B dt}{\sigma_M} = \frac{\hat{\mu}_t^R - \hat{\mu}_t^B}{\sigma_M} dt + \frac{dM_t - \hat{\mu}_t^R dt}{\sigma_M} = \frac{\hat{\mu}_t^R - \hat{\mu}_t^B}{\sigma_M} dt + d\bar{\omega}_t^R. \quad (13)$$

Investors with biased signal precision perceive $d\bar{\omega}_t^B$ to be a standard Brownian motion, since under their beliefs $\hat{\mu}_t^B$ is an unbiased estimate of μ_t . From the perspective of a rational observer, however, the rational estimate $\hat{\mu}_t^R$ will typically differ from biased investors' estimate $\hat{\mu}_t^B$; as a result, the drift term on the right-hand side of Equation (13) is (almost always) non-zero. For example, after a string of positive realizations of the signal, the biased estimate $\hat{\mu}_t^B$ is likely to exceed the rational estimate $\hat{\mu}_t^R$, resulting in a negative value for the term $\hat{\mu}_t^R - \hat{\mu}_t^B$. In such cases, biased investors will be disappointed on average by subsequent realizations of the disruption rate. The opposite predictability pattern will obtain after a string of negative realizations of the signal.

The non-zero drift of the disruption surprise factor $d\bar{\omega}_t^B$ is the source of systematic conditional return predictability in the model. Moreover, since the disruption climate is a persistent state variable, the biased investors' estimation mistake as measured by $\hat{\mu}_t^R - \hat{\mu}_t^B$ does not get corrected immediately. Therefore, the drift of the disruption surprise factor $d\bar{\omega}_t^B$ is also persistent. As we show and discuss in greater detail in Section III.B, the persistence of this drift increases the long-term variance of the factor realizations, resulting in more disperse Sharpe ratio distributions for the factor and portfolios that load on the factor.

Substituting Equation (13) in the excess return Equation (12) and taking expectations, the expected excess return on firm i from the perspective of the rational observer is

$$E^R(r_t^i) - r = \left\{ \begin{array}{l} (f_t^i V_f'(\hat{\mu}_t) + K_t^i V_K'(\hat{\mu}_t) + V_{g,z}'(\hat{\mu}_t)) (\gamma / \sigma_M^2) + \\ (\alpha \lambda K_t^i - k_z) + (k_z - \lambda K_t^i) V_K(\hat{\mu}_t) + (a_z k_z - \lambda f_t^i) V_f(\hat{\mu}_t) \end{array} \right\} \times \frac{(\hat{\mu}_t^R - \hat{\mu}_t^B)}{V_t^i}. \quad (14)$$

Cross-sectional differences in factor exposures, along with potential mispricing of the disruption surprise factor, generate cross-sectional return predictability. Specifically, some firms (i.e., early growth firms) have few projects but are accumulating new projects at a high rate relative to their current capital base. Others (i.e., no growth firms) lose their accumulated projects while not being able to replace them with new ones. Such firms, whose returns are highly sensitive to the disruption surprise factor, exhibit stronger return predictability patterns, relative to firms with low sensitivity.

II. DATA, CALIBRATION, AND SIMULATION PROCEDURES

A. Variable Definitions, Summary Statistics and Stylized Facts from the Historical Sample

The focus of this paper is on firms' valuation ratios and the underlying fundamental firm characteristics—the profitability of their assets and their growth opportunities. Before we present quantitative analyses of the model, we start by describing the measures of firm characteristics that we utilize and some stylized facts from the historical data sample.

The valuation ratio we use to capture the value anomaly is *market-to-book assets MB*, which is defined as the market value of assets (market value of equity lagged six months plus book value of debt, where the latter is computed as total liabilities plus preferred stocks minus the sum of deferred taxes and convertible debt) divided by total book assets. The empirical investment literature typically uses the book-to-market equity ratio to capture the value anomaly (e.g., Fama and French, 1992). We use *MB* because it measures the firm (not just the equity)

value and thus more closely corresponds to investors' valuations in our model, which assumes that firms are 100% equity financed. In unreported analyses we find that using book-to-market equity instead generates similar results.

The market-to-book ratio is determined in our model as a function of the growth rate and the profitability of the firm's assets. We follow Cooper, Gulen, and Schill (2008) and measure firm growth rates as *asset growth AG*, which is the percentage change in book assets from the last year to this year. To measure profitability we use *operating profitability OP*, which is defined as operating income before depreciation (total revenue minus cost of goods sold minus selling, general and administrative expenses) divided by total book assets. This measure is similar to the one used by Fama and French (2015); however, their measure of operating profits is net of interest expense and normalized by book equity. The reason for our modification is the same as explained in the previous paragraph; firms in our model are 100% equity financed and hence do not incur interest expenses. The results are again largely the same when the original Fama-French (2015) measure of operating profitability is used instead.¹⁵

Our sample includes common equity shares traded on the NYSE, AMEX and NASDAQ over the 50 year period from July 1964 to June 2014. We exclude stocks in the smallest NYSE size decile, stocks with either negative book equity or prices less than \$5 at the time of portfolio construction, and financial firms (those with one digit SIC codes of six).

For each characteristic we sort firms into five quintile portfolios. We label the top and the bottom quintile portfolios according to the underlying characteristic (e.g., the value portfolio in a

¹⁵ The literature has documented profitability-based return predictability using other measures as well. Most notably, Novy-Marx (2013) analyzes strategies based on gross profitability (i.e., revenues minus cost of goods sold divided by assets). He finds that gross profitability by itself is a somewhat weaker predictor of returns in comparison to the Fama-French (2015) measure, but that combining the value and the gross profitability strategies achieves a very high Sharpe ratio, due to the strong negative correlation between value and gross profitability.

MB sort includes stocks that are in the bottom quintile of market-to-book assets). Portfolios are formed at the end of June in each year, are value weighted, and are rebalanced monthly.¹⁶ In parts of our empirical analysis we also consider industry-adjusted characteristic portfolios, where stock returns are measured in excess of the 48 Fama-French industry portfolio returns. Monthly excess returns are computed by subtracting the one-month Treasury bill rate. Market-neutral returns are computed as the alpha plus the residual in monthly time-series CAPM regressions of the portfolio return on the value-weighted market return. While our data observations and analyses are at the monthly frequency, we report annualized volatilities and Sharpe ratios for ease of reference.

Table I reports the median values of the three firm characteristics for each characteristic sorted portfolio. As the table shows, all three characteristics exhibit substantial cross-sectional variation. Both profitability and asset growth rates are positively related to the market-to-book ratio. In particular, the firms in the high profitability portfolio resemble growth firms, with market-to-book ratios exceeding two. There is also a positive association between profitability and asset growth, but the magnitude of this relationship is quite small relative to the cross sectional variation each characteristic exhibits.

Table II reports return statistics of various investment strategies. In the first four rows of the table, we report the return standard deviations and the Sharpe ratios of four long-short portfolios: the market portfolio minus the risk-free asset, the value (i.e., low *MB*) minus the growth (i.e., high *MB*) portfolio, the high profitability minus the low profitability portfolio, and the low asset growth minus the high asset growth portfolio. The Sharpe ratio of the market

¹⁶ Portfolios are value weighted with respect to equity market capitalizations of the stocks included in them. Equally-weighted portfolios, which are dominated by smaller stocks, may generate more extreme Sharpe ratios but may not be realistic because of the implicit monthly rebalancing.

portfolio is 0.376, while the Sharpe ratios of the characteristic-sorted portfolios range from 0.281 to 0.297.

While the Sharpe ratios of the characteristic-sorted portfolios reported above are quite high, it is not clear what they should be benchmarked against, as the underlying portfolios are exposed to market risk. Because our model abstracts from a market risk factor, we primarily focus on market-neutral portfolio returns, measured as the difference between the portfolio's raw return and its estimated beta times the market excess return. Under the CAPM, the Sharpe ratios of these market-neutral portfolios are zero in expectation. Since a portfolio's realized Sharpe ratio is a scaled t -statistic, one can calculate the probability of observing a realized Sharpe ratio under the CAPM from the relevant t distribution.¹⁷ For instance, a right-tail p -value of 0.05 corresponds to a t -statistic of 1.647 with 50 years of data, which implies an annualized Sharpe ratio of 0.233 ($= 1.647 / \sqrt{50}$).

As shown in the middle panel in Table II, taking out the market exposure substantially increases the Sharpe ratios of the three characteristic-sorted strategies. The market-neutral value, asset growth, and profitability strategies attain Sharpe ratios of 0.385, 0.515, and 0.435, respectively. The corresponding p -values for these Sharpe ratios under the CAPM are 0.003, 0.0001, and 0.001. Thus, the historical Sharpe ratios of characteristic-sorted strategies have extremely low likelihoods under the null that the CAPM holds.

To see whether the documented return predictability patterns are exhibited within industries, we consider industry-adjusted (and also market-neutral) investment strategies as well, which are reported in the right-side panel of Table II. Taking out the industry returns further

¹⁷ Specifically, a portfolio's Sharpe ratio is the t -statistic on its mean return divided by \sqrt{T} , where T is the number of return observations. Our sample period is 50 years, or 600 months; therefore the relevant t distribution has 599 degrees of freedom.

increases the Sharpe ratios of the value and asset growth strategies and substantially lowers their volatilities, but has a small impact on the profitability strategy. Thus, the main return patterns of interest obtain when we sort within industries as well.

B. Model Calibration

The model has 13 parameters. Most of these parameters are difficult to calibrate based directly on data observations, and in any case, the model structure is too simplistic to provide a fully realistic description of the true data generating process. We thus pick parameters that broadly replicate the salient features of the data and highlight the mechanisms that our model is intended to capture. The calibrated parameters for our base case simulations are described in Table III. We present some comparative statics with respect to some of these parameters in Section III.D.

We set the discount rate r at 0.050. Because we think of the disruption climate as a slow moving variable, we pick a relatively small mean-reversion rate for μ , $\rho_\mu = 0.070$, which implies a half-life of shocks to μ that is approximately 10 years. The volatility of μ equals $\sigma_\mu = 0.100$. The parameter choices for ρ_μ and σ_μ imply a standard deviation of 0.267 for μ in a long time series. We pick the volatility of the transitory component of the disruption rate to be $\sigma_M = 0.250$. Thus, the long-term variation in the persistent component of the disruption rate and the short-term variation in its transitory component are similar in magnitude.

We calibrate a moderately informative signal by choosing its precision to be $\eta = 0.500$. Substituting this and the previously described parameters in Equation (9) results in $\gamma = 0.018$. Thus, in the benchmark case where rational investors perceive the correct signal precision, the

standard deviation of the investors' estimation error of μ is $\sqrt{\gamma} = 0.133$. In addition to the benchmark case of rational investors, we analyze cases with overconfident investors who perceive the signal precision to be higher than it actually is. Specifically, we set $\eta_B = 0.934$ in our base case for overconfident investors, based on survey evidence that we discuss in the next paragraph. In this case, the biased investors' perceived estimation error has a standard deviation of $\sqrt{\gamma_B} = 0.075$, which is about 56% of its rational counterpart $\sqrt{\gamma} = 0.133$.

One can evaluate the assumed degree of overconfidence in comparison to surveys that ask participants to make predictions and to report their perceived confidence intervals. In a recent study, Ben-David, Graham, and Harvey (2013) ask financial executives to project one-year S&P 500 returns and provide an 80% confidence interval. The authors find that the executives' reported confidence intervals include the realized outcome only 36.3% of the time.¹⁸ We calibrate the biased signal precision in our base case simulations to match the estimate from the Ben-David et al. Specifically, given the true signal precision $\eta = 0.5$, investors with biased precision $\eta_B = 0.934$ compute 80% confidence intervals that include the realized outcome (i.e., shocks to μ_t) 36.3% of the time on average.¹⁹

The remaining parameters of the model describe firms and their investment opportunities. We set the life expectancy of firms at 10 years, with an average of $1/q_{EG} = 3$ years spent in the early growth state, $1/q_{MG} = 4$ years spent in the mature growth state, and $1/q_{NG} = 3$ years spent

¹⁸ The standard error associated with this point estimate is 7.8%.

¹⁹ In earlier studies, Alpert and Raiffa (1969) ask Harvard Business School students to answer general knowledge questions, and Russo and Schoemaker (1992) ask money managers to answer questions about their industry. These studies respectively find the participants' 98% and 90% confidence intervals to include the correct answer 54% and 50% of the time. Our base-case overconfidence calibration implies 98% and 90% confidence intervals to include the realized outcome 60.8% and 45.5% of the time, respectively.

in the no growth state.²⁰ The average project termination rate is $\lambda = 0.150$, which implies that the average half-life of firms' active projects is 4.621 years ($= -\ln(0.5) / \lambda$).

Firms receive their projects in the early and the mature growth states. The initial investment required for a project is the same in both states, $k_{EG} = k_{MG} = k$. A higher value of k generates more cross-sectional dispersion in firms' growth rates by allowing young firms to grow faster. Accordingly, we calibrate k to match the dispersion in asset growth rates documented in Table I. Specifically, setting $k = 0.950$ approximately replicates the difference between the median asset growth rate of high AG firms and that of low AG firms, $0.483 - (-0.036) = 0.519$.

We assume that the projects firms receive in the mature growth state have half the profitability of those that they receive in the early growth state: $a_{MG} = 0.5 \times a_{EG}$. We calibrate $a_{EG} = 0.250$ (and thus $a_{MG} = 0.125$) to approximately match the median operating profitability of firms in our empirical sample, which is 0.145 in Table I. Finally, we calibrate the capital recovery rate of terminated projects $\alpha = 0.650$ to approximately match the median Tobin's q of firms in our empirical sample, which is 1.376 from Table I.

C. Simulation Procedure

Using the parameters described in the previous subsection we simulate sample paths for a set of hypothetical firms. Specifically, we start with an economy with 10,000 firms in the early growth state. As the economy evolves, these firms are endowed with new projects, their existing projects terminate, they can transition to new states, and they may die. As described in Section I.A, as firms die they are replaced with new firms born into the early growth state.

²⁰ The average number of years an individual firm appears in our empirical sample is 10.56.

The initial values of the disruption climate, μ_0 , and investors' estimate of it, $\hat{\mu}_0$, are drawn randomly from their time-invariant distributions. We simulate 200 years of data by approximating the continuous passage of time with 48 discrete time periods per year (i.e., four periods per month). We drop the first 150 years so as to allow firm characteristics to reach their steady-state distributions, and use the remaining 50 years of data (the length of our historical sample) for analysis. We repeat this procedure to generate 10,000 simulated samples.

We assign firms into value/growth portfolios based on their market-to-book ratios, asset growth portfolios based on the growth rate of their capital stock over the prior year, and profitability portfolios based on the ratio of their profitability to their capital stock. Similar to our treatment of the historical data, the cutoffs for these portfolio assignments are chosen based on quintile breakpoints of the underlying characteristics.

III. QUANTITATIVE ANALYSIS OF THE MODEL

A. Summary Statistics of Simulated Samples

Table IV presents summary statistics of the simulated data samples generated from the base case calibration where the investors' perception of signal quality is biased. While some of the reported statistics, such as capital stock, profitability, and asset growth do not depend on the perceived or true signal precision in any case, others, such as valuation ratios and returns, are affected by the signal precision. However, in unreported analyses we find that the summary statistics reported in Table IV exhibit negligible sensitivity in this regard.

As shown in Panel A, the market portfolio, defined as the value-weighted portfolio of all simulated firms, earns an average annual return of 5.00%, matching the discount rate r . The median firm's capital stock is 3.150, profitability is 0.144, and Tobin's q is 1.371. Recall that the

model is calibrated to approximately match the latter two statistics. Importantly, none of the statistics reported in Panel A exhibit substantial variation across samples or across years within samples, indicating that the creative destruction process in the model does not materially affect market-wide aggregates.

Panel B reports the time-series averages of the median values of firm fundamentals in each of the three growth states. The mature growth firms are the largest and the early growth firms are the smallest on average. Profitability is the highest for firms in the early growth state. Since firms do not receive any projects following their transitions from mature growth to no growth, and since these transitions occur with equal likelihood for any firm in the mature growth state, by construction the mature growth and the no growth states exhibit the same average profitability. Asset growth rates decline as firms transition from early growth to mature growth, and become negative in the no growth state. Similarly, Tobin's q values are the highest in the early growth state, followed by the mature growth state and then the no growth state.

Panel C reports statistics on characteristic-sorted portfolios. The first three columns report the percentages of firms in a given characteristic-sorted portfolio that belong to each of the three growth states. The next four columns are the median values of firm fundamentals for each portfolio. The growth portfolio (i.e., the high market-to-book portfolio) consists mainly of firms in the early growth state, whereas the value portfolio (i.e., the low market-to-book portfolio) consists mainly of firms in the no growth state. Growth firms also tend to be smaller, more profitable, and exhibit higher growth rates relative to value firms. High asset growth firms, which are primarily in the early growth state, are larger and have higher Tobin's q values. However, there is little difference between the profitability rates of high versus low asset growth firms. Profitability portfolios are somewhat more evenly distributed across the three growth

states. High profitability firms are larger and exhibit higher Tobin's q values; however, in contrast to the empirical sample, they exhibit slightly lower growth rates than low profitability firms. Overall, the cross-sectional characteristics of the model-simulated firms appear to be broadly in line with the salient features of the actual data.

Panel D reports the annualized return standard deviations and correlations of the three characteristic-sorted strategies. The simulated strategies are less volatile than their empirical counterparts in Table I, and exhibit strong positive correlations with each other. This disconnect between the simulated and historical data is not surprising. Our model, which focuses only on the disruptive innovation factors, abstracts from industry and other shocks that may apply to certain groups of firms but not others (e.g., shocks to consumer demand that generate common earnings movements among consumer goods producers). Such risk factors are likely to increase the return volatilities of the characteristic-sorted strategies and dampen the correlations between their returns.²¹ Of course, if extraneous risk factors add to volatilities of characteristic-sorted strategies, without contributing to mean returns, then strategies that explicitly filter out these risks should exhibit higher Sharpe ratios than the relatively simple characteristic-sorted strategies we consider. The generally higher Sharpe ratios that obtain with industry-adjusted strategies in Table I is consistent with this view.

B. Sharpe Ratio Distributions

We are primarily interested in the model-generated Sharpe ratios of various portfolios, which we report in Table V and in Figures 1 and 2. As a first step, we consider the Sharpe ratio

²¹ A second reason why the model generates relatively low volatilities is that it abstracts from financial leverage effects. In unreported analyses we find that adjusting for financial leverage in historical data reduces volatilities of the characteristic-sorted strategies by up to 30%. Importantly, Sharpe ratios are invariant to leverage-induced changes in volatility, so our Sharpe ratio analyses should be relevant despite the model's abstractions in this regard.

of a factor portfolio that is constructed to be perfectly correlated with the model's potentially mispriced factor, i.e., the disruption surprise $d\bar{\omega}_t$ (see Equations 7 and 13). A useful benchmark for the Sharpe ratios generated by our model is the distribution of the Sharpe ratios associated with this factor portfolio when investors are rational, i.e., when $\eta_B = \eta$. When this is the case, the disruption surprise factor $d\bar{\omega}_t$ is a standard Brownian motion by construction, and the cumulative abnormal return (CAR) of the factor portfolio over a time period Δt is distributed normally with mean zero and standard deviation $\sqrt{\Delta t}$. Moreover, any investment strategy with constant exposure to this factor portfolio for T years (and zero exposure to the soft signal factor) can be shown to have an annual Sharpe ratio that is also distributed normally with mean zero and standard deviation $\sqrt{1/T}$, which equals $\sqrt{1/50} = 0.141$ for our 50-year sample period.²²

The black dotted curve in Figure 1 plots this benchmark distribution, and Panel A of Table V provides the z -values and the Sharpe ratios that correspond to various percentiles of the distribution.²³ This theoretical benchmark can be compared to the realized distribution of the factor portfolio's Sharpe ratios generated in our simulations with rational investors, as shown in the first row of Panel B. As can be seen by comparing Panels A and B, the simulated distribution closely overlaps with the theoretical benchmark.²⁴

Figure 1 plots three other Sharpe ratio distributions, those of the value, the asset growth, and the profitability portfolios, that are generated in simulations with rational investors.²⁵ Panel

²² Note that the distribution of the Sharpe ratio collapses to zero as the length of the sample period T becomes arbitrarily large.

²³ The z -values, which are standardized Sharpe ratios (i.e., Sharpe ratio divided by 0.141, the standard deviation of the distribution) are provided for ease of reference. Return predictability is assessed in most empirical work with magnitudes of t -statistics, which converge to z -values in large samples.

²⁴ The very small differences are due to simulation sampling and numerical approximation errors.

²⁵ For visual ease Figures 1 and 2 plot smoothed probability distribution function estimates that are obtained by applying a normal kernel function to simulated Sharpe ratios.

B of Table V presents the percentile values for those distributions. The distributions of the characteristic portfolios have slightly negative means and are also left-skewed, indicating that these investment strategies are somewhat more likely to generate negative Sharpe ratios compared to the benchmark in Panel A. The deviations from the benchmark are likely to result in part from numerical approximation errors in solving for model firm valuations. However, the deviations also reflect the fact that the returns on characteristic portfolios, in contrast to the factor portfolio, are not distributed i.i.d. normal.²⁶ Importantly, the deviations from the benchmark Sharpe ratio distribution are both small in economic magnitude, and go in the opposite direction of the historically observed Sharpe ratio patterns.

The main results of our analysis, the distributions of Sharpe ratios under alternative assumptions about investor beliefs, are described in Figure 2 and Panel C of Table V. Figure 2 plots the Sharpe ratio distributions for the value minus growth portfolio with unbiased and biased beliefs about the signal precision.²⁷ Panel C of Table V replicates the analysis in Panel B under biased beliefs. As the figure and the table show, introducing the overconfidence bias results in more dispersed Sharpe ratio distributions relative to the benchmark case with unbiased beliefs. The economic magnitudes of the Sharpe ratios in the distributions' tails are large. For instance, at the 90th percentile, the disruption surprise factor has a Sharpe ratio of 0.424, while the Sharpe ratios of the value, asset growth, and profitability portfolios are 0.367, 0.326, and 0.332, respectively.

Our simulated Sharpe ratios can be compared to the historical Sharpe ratios of the

²⁶ Indeed, in unreported analyses we find that the realized mean returns of the value, the asset growth, and the profitability portfolios are positively correlated with the realized time-series return standard deviation of those portfolios (correlations of 0.64 for the value, 0.34 for the asset growth, and 0.55 for the profitability strategies across the 10,000 sample paths). Thus, in sample paths where the characteristic portfolios earn positive returns on average, those returns turn out to be more volatile, dampening the realized Sharpe ratio. Such deviations from constant return volatility are not surprising in a model like ours, where the factor affects firms' valuations in non-linear ways.

²⁷ Figures for the asset growth and the profitability portfolios are omitted for brevity.

market-neutral portfolios reported in Panel B of Table I. Specifically, the market-neutral value, asset growth, and profitability portfolios have historical Sharpe ratios of 0.385, 0.515, and 0.435, respectively. As we discussed earlier, the probability of observing Sharpe ratios that equal or exceed these magnitudes under the null hypothesis that the expected returns are always zero is exceedingly small (p -value of 0.003 for the value portfolio and even smaller for the asset growth and profitability portfolios). In contrast, the probability of exceeding the historical Sharpe ratios in our simulated model with biased investors is 0.089 for the value portfolio, 0.026 for the asset growth portfolio, and 0.046 for the profitability portfolio. In other words, Sharpe ratios that match historical magnitudes are somewhat unlikely in our model, but they are significantly more likely than under the traditional null.

It should be stressed that these probabilities are generated in a model where there is no unconditional return predictability, so the higher probability of extreme outcomes arises because biased beliefs increase the dispersion of the Sharpe ratios. It is relatively straightforward to understand why biased beliefs can increase the dispersion of the factor portfolio's Sharpe ratio. As shown in Equation (13), the factor has a persistent drift that arises because of the accumulation of investors' past mistakes in interpreting the soft signal. This persistent drift generates positive serial correlation in factor portfolio returns. Because of this positive serial correlation, the short-term standard deviations that are used to calculate the Sharpe ratios are too low relative to the long-term standard deviations that determine the dispersion of these Sharpe ratios.

To illustrate these patterns, the first column of Panel A of Table VI reports serial correlations of the factor portfolio returns. We calculate these serial correlations with returns measured at one-month, one-year, two-year, and five-year intervals in each simulated sample.

The table reports the means and the standard deviations computed across 10,000 samples, as well as the fraction of samples in which the estimated serial correlation is positive. As the table shows, factor returns measured over various intervals exhibit positive serial correlation on average and in majority of the simulated samples.²⁸ In Panel B of Table VI we report variance ratios, defined as the annualized variance of returns measured over a long period (e.g., five years) divided by the annualized variance of monthly returns. The variance ratios tend to exceed one, which is consistent with the positive serial correlation, and they increase as longer-term returns are considered. For instance, five-year returns are about 40% more volatile on an annualized basis than monthly returns ($\sqrt{1.958} = 1.40$).

The characteristic-sorted portfolios are exposed to the disruption surprise factor, which exhibits positive serial correlation, but they are also exposed to the other systematic risk factor, the soft information signal. Although the soft information signal is not serially correlated, realizations of this signal cause mispricing, which is later reversed. The build-up and the subsequent correction of mispricing can thus generate a source of negative serial correlation that attenuates the positive serial correlation induced by the disruption surprise factor.

The last three columns of Table VI, which report serial correlations and variance ratios for the three characteristic-sorted portfolios, indicate that this is indeed the case. As the table shows, the portfolios do exhibit positive serial correlation and high long-term volatility relative to short-term volatility. The patterns of both statistics are similar to those for the disruption surprise factor, but are somewhat attenuated. This attenuation effect is consistent with the characteristic portfolios being subject to the second source of systematic risk, the soft

²⁸ Monthly returns exhibit weak positive serial correlation, which is not surprising, since short-term return variation is determined by the diffusion term rather than the drift term.

information signal.

It is of interest to compare the serial correlations and variance ratios generated by our model to the historical evidence. The serial correlations of the historical annual returns of the value, the asset growth, and the profitability portfolios are -0.055 , 0.098 , and 0.147 , respectively (not reported in tables). These figures are within the range of annual serial correlations from the model reported in Table VI. For instance, while the average serial correlation for the annual returns of the value portfolio is 0.129 in the model, about 20% of the simulated samples exhibit *negative* serial correlation ($1 - 0.799 = 0.201$). As for long-term volatility, the variance ratios of historical annual to monthly returns of the value, the asset growth, and the profitability portfolios are 1.383 , 1.237 , and 1.586 , respectively. These figures, which are in fact somewhat higher than their simulated counterparts, suggest that long-term returns are substantially more volatile than short-term returns in the historical data as well.

C. Implementable Investment Strategies

Up to this point we have explored the returns of static characteristic-sorted portfolios. In this section we examine the Sharpe ratios that are achieved by dynamic strategies that may be conditioned on information that is revealed along the simulated sample paths. Our interpretation is that we are examining the success of “quant” investors who are somewhat more rational, in a way that we explicitly describe below, than the overconfident investors that set prices in our model. The expected success of these quant investors provides a way to gauge the magnitude of the inefficiency of our simulated stock market, and as such, helps us interpret the results from the last section.

Our analysis of quant investors requires a number of assumptions. First, we assume that

the quant is risk averse with preferences described by log utility. Risk aversion is necessary to prevent the investor from taking infinite positions in mispriced assets and log utility simplifies the portfolio choice problem by shutting down hedging demands. Let $\{\bar{r}_t, \sigma_t\}$ describe the investor's opportunity set at time t , where \bar{r}_t and σ_t are the risky asset's conditional expected excess return and standard deviation.²⁹ The weight a log utility investor assigns to the risky asset is given by

$$\omega_t = \frac{\bar{r}_t}{\sigma_t^2}. \quad (15)$$

Since asset prices in our model are set by a risk-neutral group of agents with homogenous beliefs, the quant investor's trades do not influence prices. In other words, the mass of risk neutral investors are willing to trade any quantity of any asset with the risk-averse quant without revising their beliefs—i.e., they and the quant investor would agree to disagree—and the resulting trades would have no price impact.³⁰

We consider three types of quant investors. The first quant is a fully rational investor who employs “the optimal factor timing strategy.” Specifically, the fully rational investor observes both the realized disruption rate dM_t and the soft information signal ds_t and in addition knows the correct signal precision. Thus, the fully rational quant can infer the biases of the other investors and can thus compute the conditional drift of the disruption surprise factor

²⁹ We consider quant investors who invest in a single risky asset. A generalization to multiple risky assets is straightforward, but not particularly relevant in our calibrated model, since risky assets in our simulations exhibit strong correlations.

³⁰ One can envision more general models where all agents are risk averse and quant investors have non-zero measure, generating price impact for their trades. Given our partial-equilibrium focus on a quant investor and the profitability of her trading strategies, such a general model is beyond the scope of the current paper.

$d\bar{\omega}_t^B$. Using Equation (13), the investment opportunity set for this investor is characterized by

$$\{\bar{r}_t, \sigma_t\} = \left\{ \frac{\hat{\mu}_t^R - \hat{\mu}_t^B}{\sigma_M}, 1 \right\}. \quad (16)$$

The optimal weight on the disruption surprise factor $d\bar{\omega}_t^B$ is then given by Equation (15).

Inspecting Equations (15) and (16), the optimal factor strategy attempts to time the mispriced factor by increasing exposure to it when the rational estimate $\hat{\mu}_t^R$ exceeds the biased estimate $\hat{\mu}_t^B$ by a greater amount.

Against this benchmark of a fully rational quant, we consider two ad hoc, but probably more realistic depictions of quant investors. The first ad hoc quant completely ignores the soft information signal and updates his estimate of μ_t using only the realized disruption rate dM_t . The motivation for considering this case is that many of the quantitative investment strategies observed in practice are designed to utilize only hard information. Specifically, the quant investor in this case updates his estimate $\hat{\mu}_t^Q$ using the formulas in Equations (6) through (9) but by setting $\eta = 0$, and chooses his exposure to the disruption surprise factor $d\bar{\omega}_t^B$ using Equation (15) based on his perceived investment opportunity set

$$\{\bar{r}_t, \sigma_t\} = \left\{ \frac{\hat{\mu}_t^Q - \hat{\mu}_t^B}{\sigma_M}, 1 \right\}. \quad (17)$$

The second ad hoc quant employs a “characteristic timing strategy” that bases trades on

the past returns of characteristic-sorted long-short portfolios. The “characteristic timing quant” estimates the expected excess return \bar{r}_t and the return standard deviation σ_t of a characteristic-sorted portfolio using the past 10, 25, or 50 years of return data, and then chooses his portfolio weights by using these quantities according to Equation (15). These heuristic investment strategies closely resemble real-world quantitative approaches, and since they are purely data driven, they do not require the quant investor to know the model structure.

Table VII summarizes the results for the optimal factor timing strategy, the quantitative factor timing strategy, and the characteristic timing strategies. For each strategy, we compute the investor’s annualized Sharpe ratio over each simulated 50-year sample period and report the distributions of these Sharpe ratios across all sample paths (analogous to Table V).

Panel A contains the results for the two factor strategies, the optimal factor timing strategy that conditions on both hard and soft information, and the quantitative factor timing strategy that utilizes only the hard information gleaned from realizations of the disruption rate. The optimal strategy attains a Sharpe ratio of 0.413 in the median sample path. Interestingly, the quantitative factor timing strategy, which does not condition on the soft information signal at all, does almost as well as the optimal strategy. This indicates that future realizations of hard information are an effective substitute for the current realization of soft information in this calibration, so that ignoring the latter does not cause a material deterioration in investment performance.

Panel B reports the results for the characteristic strategies. As indicated above, these strategies may provide a more realistic perspective on quant investors’ ability to detect mispricing: unlike the factor strategies that rely on optimal filtering rules, the characteristic strategies do not require the quant investor to have any knowledge of the model parameters. The

Sharpe ratio distributions of the characteristic strategies show that using more recent data for estimation achieves the strongest performance. For instance, the value timing strategy achieves a median Sharpe ratio of 0.216 when the quant investor forms his portfolio using data from the past 10 years, but a median Sharpe ratio of only 0.058 when data from the past 50 years is used. Even with a 10-year estimation period, however, there is more than 10% chance that the quant investor obtains a negative Sharpe ratio.

As we mentioned at the beginning of this section, the Sharpe ratios of the three types of timing strategies we consider provide one way to gauge the efficiency of our model's stock market. If, for example, the model generates high Sharpe ratios for characteristic-sorted portfolios in about 10% of the sample paths but low average Sharpe ratios (e.g., less than 0.10) for quant investors, then one might conclude that the historical return patterns can be generated in a market that is only mildly inefficient. The results of this section, however, suggest that the historical evidence does in fact require fairly substantial market inefficiency.

D. Comparative Statics and Alternative Specifications

We conclude the numerical analysis of the model by considering comparative statics and alternative specifications. The goal of this exercise is to provide intuition about what is driving our results and to better understand the conditions under which the Sharpe ratios observed over the past 50 years can be generated.

Table VIII, which summarizes this analysis, reports for brevity only a subset of the Sharpe ratio statistics considered in earlier subsections. Specifically, we focus on the 90th and the 95th percentiles (i.e., the right tail) of the Sharpe ratio distribution of the value minus growth portfolio. We also report the 50th and the 90th percentiles of the Sharpe ratios attained by the

optimal factor timing strategy and the value timing strategy that utilizes past 10 years of data.

Panel A contains the base-case calibration values of previously reported statistics for comparison. Panel B contains the comparative statics with respect to the degree of overconfidence, where we keep the true signal precision η constant and vary the perceived signal precision η_B . Recall that our base-case calibration of overconfidence is based on the survey evidence in Ben-David, Graham, and Harvey (2013), which indicates that survey respondents' 80% confidence intervals contain the realized outcome only 36.3% of the time. The cases of low and high overconfidence in Panel B correspond to 80% confidence intervals that contain the realized outcome only 50% and 20% of the time, respectively. The table also reports the case of maximum overconfidence, where investors perceive their soft signal to be infinitely precise and thus ignore the information in realized disruption rates in forming their expectations. This case is clearly unrealistic, but it is informative about the bounds on return predictability the model can generate.

These simulations indicate that the magnitudes of the historical Sharpe ratios are very unlikely to occur when overconfidence is low, but they do occur in more than 10% of the sample paths when overconfidence is high. As expected, the quant strategies are also more effective when investors are more overconfident. Our simulations indicate that when overconfidence is high, the fully rational quant that implements the optimal factor timing strategy generates extremely high Sharpe ratios on average, while the quant that implements the value timing strategy is also quite high. In other words, the degree of overconfidence required to generate high Sharpe ratios for characteristic-sorted portfolios in the tails generates high *expected* Sharpe ratios for quant investors.

Another aspect of the model that contributes to return predictability in relatively long

samples is the persistence of the disruption climate. If disruption prospects in the economy change very quickly, investors' mistakes across different years would be largely uncorrelated, i.e., some years they will be overly optimistic about growth firms and other years they will be overly optimistic about value firms. When this is the case, characteristic-sorted portfolios are not likely to exhibit significant return predictability over a 50-year sample period. The base-case calibration assumes $\rho_\mu = 0.07$, which corresponds to a half-life of disruption shocks that is approximately 10 years. The question we now ask is whether our results change significantly if we assume the half-life of the disruption shocks to be somewhat longer or shorter.

In Panel C of Table VIII, we consider three alternative specifications: (i) a case of very low persistence, where $\rho_\mu = 0.693$ implies a half-life of one year for disruption shocks; (ii) a case of relatively low persistence, where $\rho_\mu = 0.139$ implies a half-life of five years; (iii) a case of relatively high persistence, where $\rho_\mu = 0.046$ implies a half-life of 15 years. As the table shows, biased beliefs do not materially affect the dispersion of Shape ratio distributions when disruption shocks die off quickly, i.e., with a half-life of one year. However, with plausible levels of persistence, i.e., when disruption shocks have a half-life of five years or longer, Sharpe ratio distributions exhibit substantial dispersion. The results also show that, with a 50-year sample period, increasing persistence beyond the base-case calibration (i.e., an increase in half-life from 10 to 15 years) has negligible impact on Sharpe ratio distributions.

Up to now, the only source of return predictability we consider is investors' overconfidence about the precision of their soft information. We now briefly consider an additional source of bias, namely, investors being initially too optimistic about the disruption climate. As Shiller (2000) points out, technological innovations such as the internet tend to create expectations of substantial opportunities and changes in the business landscape, even though the

immediate commercial potential of the new technology may be far from clear. In our model, investors' expectations may be initially biased in the manner suggested by Shiller; however, over time they do learn, so in the steady-state, the return predictability is not influenced by this bias. Our implicit assumption is that in 1964, the start of our sample, because of data and computational limitations, investors had not learned a lot from their past history and that investor beliefs were influenced by this inherent optimistic bias.

The way we introduce initial investor optimism is as follows. At the beginning of each simulated 50-year sample period, we allow for an exogenous, ad-hoc increase in investors' estimate of the disruption climate, $\hat{\mu}_t$.³¹ The rest of the model applies the same as before: investors receive new signals, observe the realized disruption rates, and update their beliefs. We implement the initial optimism experiment in two different specifications. The first is the fully rational case where investors have unbiased beliefs about the signal precision η . In this case, investors process the new information rationally, but start the sample period with optimistic beliefs. The second specification is our-base case calibration where investors are overconfident about the precision of their signal. Since overconfident investors update their beliefs more slowly in response to the realized disruption rate, the effects of initial optimism in this second specification are likely to be more persistent.

Panel D of Table VIII summarizes the results. When investors are unbiased about their signal precision, initial optimism generates very weak return predictability. For instance, the tail Sharpe ratios of the value minus growth portfolio are only slightly higher than their counterparts in the fully rational case in Table V. Thus, the effects of initial optimism quickly dissipate when investors update their priors rationally. Sharpe ratios become substantially larger when initial

³¹ The specific magnitude we add to $\hat{\mu}_t$ is two times the standard deviation of the time-invariant distribution of μ .

optimism is added to the model with overconfident investors. In this case, the initial bias caused by optimism takes longer to correct since investors are also fairly confident about the precision of their initial beliefs.

The main takeaway from our comparative statics is that if investors are not initially overly optimistic about the disruption climate then the observed Sharpe ratios either reflect extremely overconfident investors or are the realization of a very unusual sample path. However, a combination of initial optimism and less extreme overconfidence can generate the historical Sharpe ratios of characteristic-sorted portfolios in more than 10% of the simulated sample paths. It is also the case that the characteristic-timing strategies generate particularly high Sharpe ratios in this setting, which is indicative of a relatively inefficient market.

IV. CONCLUSION

Over the past 50 years, market-neutral portfolios formed on characteristics like value, profitability and asset growth have generated extremely high Sharpe ratios. Whether these return patterns reflect positive unconditional expected returns, and whether they are likely to repeat in the future, are questions of interest to financial economists and investment professionals alike. Conventional asset pricing tests, which are designed to reject the null hypothesis that expected excess returns are zero at all times, are less useful in answering these questions. In particular, as we stress in this paper, if conditional expected returns can vary, the unconditional expected rates of return are estimated much less precisely. This implies that the past performance of a characteristic-sorted portfolio may not be a good indicator of its future returns, even if the structure of the economy remains the same.

To formalize these points we develop and calibrate a dynamic model where characteristic-sorted portfolio returns are linked to a systematic factor that determines the cross-section of firm fundamentals such as profitability and growth. The unconditional expected return associated with this factor is zero, but over finite intervals the factor generates persistent excess returns due to conditional mispricing. As simulations of the calibrated model illustrate, the persistence generated by the model greatly increases the probability of observing high Sharpe ratios for characteristic-sorted portfolio returns in samples of comparable length to the historical sample. Overall, our results suggest that the historical evidence on return predictability may provide a poor guide in assessing future expected returns.

For the sake of parsimony we have made a number of assumptions that allow us to focus on the persistence of characteristic-sorted portfolios returns rather than their unconditional means. In particular, we assume risk neutrality and design a model where one economic concept generates multiple anomalies. Given our simplifications, it is not surprising that our model does not capture all salient features of the data. For instance, the model generates characteristic-sorted portfolio return variances that are too low and correlations that are too high. These discrepancies can potentially be addressed by incorporating insights from our model into existing risk-based models that generate a richer structure of expected returns.

Finally, while our main focus is on overconfidence as the source of mispricing, there are likely to be a number of other impediments to learning that may have influenced historical stock returns. For example, although investors in our model are overconfident about the precision of their soft signal, they otherwise use Bayes' rule to process all available information. In reality, there are obvious information processing costs, especially in the early part of our sample, which clearly affected learning. Indeed, financial economists were also slow to learn about return

anomalies, even after having access to large data bases and fast computers. So, slow learning in reality may have as much to do with technical or institutional impediments as well as with biased perceptions. Future research can directly model those impediments, along with the innovations that relax them and make markets more efficient over time.

APPENDIX A. DERIVATIONS OF FIRM VALUATIONS AND RETURNS

In this Appendix we provide the details of the derivations of firm value and return equations. As shown in Equation (10), the value of firm i equals the discounted value of its expected cash flows conditional on all available information:

$$V_t^i = E_t \left[\int_{u=t}^{\infty} e^{-r(u-t)} \left[f_u^i du + (\alpha \lambda K_u^i - k_z)(du + dM_u) \right] \right]. \quad (\text{A.1})$$

Using Equation (7), we write

$$dM_u = \hat{\mu}_u du + \sigma_M d\bar{\omega}_u. \quad (\text{A.2})$$

Substituting in Equation (A.1), taking the expectation with respect to the Brownian term $d\bar{\omega}_u$, and writing the firm value as a function of the state variables, we have

$$V(z_t^i, K_t^i, f_t^i, \hat{\mu}_t) \equiv E_t \left[\int_{u=t}^{\infty} e^{-r(u-t)} \left[f_u^i + (\alpha \lambda K_u^i - k_z)(1 + \hat{\mu}_u) \right] du \right]. \quad (\text{A.3})$$

The first state variable z_t^i is the firm's idiosyncratic growth state, which follows the Markov process described in Section II.A. The laws of motion of the other three state variables are

$$dK_t^i = (k_z - \lambda K_t^i) \left((1 + \hat{\mu}_t) dt + \sigma_M d\bar{\omega}_t \right), \quad (\text{A.4})$$

$$df_t^i = (a_z k_z - \lambda f_t^i) \left((1 + \hat{\mu}_t) dt + \sigma_M d\bar{\omega}_t \right), \quad (\text{A.5})$$

$$d\hat{\mu}_t = -\rho_\mu \hat{\mu}_t dt + \sigma_\mu \eta ds_t + \frac{\gamma}{\sigma_M} d\bar{\omega}_t. \quad (\text{A.6})$$

Equations (A.4) and (A.5) follow from substituting Equation (A.2) in Equations (1) and (2), respectively. Equation (A.6) restates Equation (6).

Using Itô's Lemma, the instantaneous rate of return of the firm dr_t^i is given by

$$\begin{aligned} dr_t^i V_z &= f_t^i dt + (\alpha \lambda K_t^i - k_z) \left((1 + \hat{\mu}_t) dt + \sigma_M d\bar{\omega}_t \right) + dQ_{z,z+1} (V_{z+1} - V_z) \\ &+ (k_z - \lambda K_t^i) \left((1 + \hat{\mu}_t) dt + \sigma_M d\bar{\omega}_t \right) \frac{\partial V_z}{\partial K} + (a_z k_z - \lambda f_t^i) \left((1 + \hat{\mu}_t) dt + \sigma_M d\bar{\omega}_t \right) \frac{\partial V_z}{\partial f} \\ &+ \left(-\rho_\mu \hat{\mu}_t dt + \sigma_\mu \eta ds_t + \frac{\gamma}{\sigma_M} d\bar{\omega}_t \right) \frac{\partial V_z}{\partial \mu} \\ &+ \frac{(k_z - \lambda K_t^i)^2 \sigma_M^2}{2} \frac{\partial^2 V_z}{\partial K^2} + \frac{(a_z k_z - \lambda f_t^i)^2 \sigma_M^2}{2} \frac{\partial^2 V_z}{\partial f^2} + \left(\frac{\sigma_\mu^2 \eta^2}{2} + \frac{\gamma^2}{2\sigma_M^2} \right) \frac{\partial^2 V_z}{\partial \mu^2} \\ &+ (k_z - \lambda K_t^i) (a_z k_z - \lambda f_t^i) \sigma_M^2 \frac{\partial^2 V_z}{\partial K \partial f} + (k_z - \lambda K_t^i) \gamma \frac{\partial^2 V_z}{\partial K \partial \mu} + (a_z k_z - \lambda f_t^i) \gamma \frac{\partial^2 V_z}{\partial f \partial \mu}. \end{aligned} \quad (\text{A.7})$$

In Equation (A.7), the term V_z indicates the firm value when the firm is in growth state

$z \in \{EG, MG, NG\}$. Note that this notation suppresses the state variables $\{K_t^i, f_t^i, \hat{\mu}_t\}$. With slight

abuse of notation we write V_{z+1} to indicate the firm value in the next growth state the firm will

transition into (e.g., $z + 1 = MG$ when $z = EG$). Similarly, we use the notation $dQ_{z,z+1}$ for the

Poisson process governing the growth state, and later below we use $q_{z,z+1}$ to denote the transition

rate (e.g., $q_{z,z+1} = q_{EG}$ when $z = EG$). When the firm is in the no growth state $z = NG$, V_{z+1} corresponds to the firm's liquidation value (i.e., the sum of the present value of active projects' cash flows and the value expected from recovery of capital), which we characterize below.

Computing the expected value of the right-hand side of Equation (A.7) and using the fact that the firm's expected return equals the discount rate r , we obtain the partial differential equation that characterizes the firm value:

$$\begin{aligned}
rV_z = & f_t^i + (\alpha\lambda K_t^i - k_z)(1 + \hat{\mu}_t) + q_{z,z+1}(V_{z+1} - V_z) \\
& + (k_z - \lambda K_t^i)(1 + \hat{\mu}_t) \frac{\partial V_z}{\partial K} + (a_z k_z - \lambda f_t^i)(1 + \hat{\mu}_t) \frac{\partial V_z}{\partial f} - \rho_\mu \hat{\mu}_t \frac{\partial V_z}{\partial \mu} \\
& + \frac{(k_z - \lambda K_t^i)^2 \sigma_M^2}{2} \frac{\partial^2 V_z}{\partial K^2} + \frac{(a_z k_z - \lambda f_t^i)^2 \sigma_M^2}{2} \frac{\partial^2 V_z}{\partial f^2} + \frac{1}{2} \left(\sigma_\mu^2 \eta^2 + \frac{\gamma^2}{\sigma_M^2} \right) \frac{\partial^2 V_z}{\partial \mu^2} \\
& + (k_z - \lambda K_t^i)(a_z k_z - \lambda f_t^i) \sigma_M^2 \frac{\partial^2 V_z}{\partial K \partial f} + (k_z - \lambda K_t^i) \gamma \frac{\partial^2 V_z}{\partial K \partial \mu} + (a_z k_z - \lambda f_t^i) \gamma \frac{\partial^2 V_z}{\partial f \partial \mu}.
\end{aligned} \tag{A.8}$$

Equation (12) in the text results from subtracting Equation (A.8) (with the dt terms included) from Equation (A.7), rearranging terms, and defining the firm's idiosyncratic return as

$$d\varepsilon_t^i = (dQ_{z,z+1} - q_{z,z+1} dt) \frac{V_{z+1} - V_z}{V_z}. \tag{A.9}$$

We now conjecture that the firm value in state z is separable with the following functional form:

$$V_z = f_t^i V_f(\hat{\mu}_t) + K_t^i V_K(\hat{\mu}_t) + V_{g,z}(\hat{\mu}_t). \quad (\text{A.10})$$

Substituting the conjectured functional form into Equation (A.8), we obtain

$$(r + \lambda(1 + \hat{\mu}_t))V_f + (\gamma\lambda + \rho_\mu \hat{\mu}_t)V_f' - \frac{1}{2} \left(\sigma_\mu^2 \eta^2 + \frac{\gamma^2}{\sigma_M^2} \right) V_f'' = 1, \quad (\text{A.11})$$

$$(r + \lambda(1 + \hat{\mu}_t))V_K + (\gamma\lambda + \rho_\mu \hat{\mu}_t)V_K' - \frac{1}{2} \left(\sigma_\mu^2 \eta^2 + \frac{\gamma^2}{\sigma_M^2} \right) V_K'' = \alpha\lambda(1 + \hat{\mu}_t), \quad (\text{A.12})$$

$$\begin{aligned} rV_{g,z} + \rho_\mu \hat{\mu}_t V_{g,z}' - \frac{1}{2} \left(\sigma_\mu^2 \eta^2 + \frac{\gamma^2}{\sigma_M^2} \right) V_{g,z}'' - q_{z,z+1} (V_{g,z+1} - V_{g,z}) = \\ k_z (1 + \hat{\mu}_t) (V_K - 1) + k_z \gamma V_K' + (a_z k_z) ((1 + \hat{\mu}_t) V_f + \gamma V_f'). \end{aligned} \quad (\text{A.13})$$

Equations (A.11) and (A.12) are ordinary differential equations characterizing the functions

$V_f(\hat{\mu}_t)$ and $V_K(\hat{\mu}_t)$, respectively. Equation (A.13) is a system of ordinary differential equations

characterizing the solutions of the functions $V_{g,z}(\hat{\mu}_t)$. Note that $V_{g,z}(\hat{\mu}_t) = V_{g,z+1}(\hat{\mu}_t) = 0$ when

the firm is in the no growth state $z = NG$. We solve the ordinary differential equations in

Equation (A.11) through (A.13) numerically using Chebyshev polynomial approximations.

REFERENCES

- Alpert, Marc, and Howard Raiffa, 1969, A progress report on the training of probability assessors, unpublished manuscript, Harvard University.
- Alti, Aydogan, and Paul C. Tetlock, 2014, Biased beliefs, asset prices, and investment: A structural approach, *Journal of Finance* 69, 325-361.
- Baker, Malcolm, and Jeffrey Wurgler, 2006, Investor sentiment and the cross-section of stock returns, *Journal of Finance* 61, 1645-1680.
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 307-343.
- Ben-David, Itzhak, John R. Graham, and Campbell R. Harvey, 2013, Managerial miscalibration, *Quarterly Journal of Economics* 128, 1547-1584.
- Berk, Jonathan B., Richard C. Green, and Vasant Naik, 1999, Optimal investment, growth options, and security returns, *Journal of Finance* 54, 1553-1607.
- Carlson, Murray, Adlai Fisher, and Ron Giammarino, 2004, Corporate investment and asset price dynamics: Implications for the cross-section of returns, *Journal of Finance* 59, 2577-2603.
- Cooper, Michael J., Huseyin Gulen, and Michael J. Schill, 2008, Asset growth and the cross section of stock returns, *Journal of Finance* 63, 1609-1651.
- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and overreactions, *Journal of Finance* 53, 1839-1886.
- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 2001, Overconfidence, arbitrage, and equilibrium asset pricing, *Journal of Finance* 56, 921-965.
- Fama, Eugene F., and Kenneth R. French, 1992, The cross-section of expected stock returns, *Journal of Finance* 47, 427-465.
- Fama, Eugene F., and Kenneth R. French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1-22.
- Gârleanu, Nicolae, Leonid Kogan, and Stavros Panageas, 2012, Displacement risk and asset returns, *Journal of Financial Economics* 105, 491-510.
- 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.

- Kogan, Leonid, Dimitris Papanikolaou, and Noah Stoffman, 2015, Winners and losers: Creative destruction and the stock market, working paper.
- Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2015, Interpreting factor models, working paper.
- Lewellen, Jonathan, and Jay Shanken, 2002, Learning, asset-pricing tests, and market efficiency, *Journal of Finance* 57, 1113-1145.
- MacKinlay, A. Craig, 1995, Multifactor models do not explain deviations from the CAPM, *Journal of Financial Economics* 38, 3-28.
- McLean, R. David, and Jeffrey Pontiff, 2016, Does academic research destroy stock return predictability? forthcoming, *Journal of Finance*
- Novy-Marx, Robert, 2013, The other side of value: The gross profitability premium, *Journal of Financial Economics* 108, 1-28.
- Pastor, Lubos, and Robert F. Stambaugh, 2012, Are stocks really less volatile in the long run? *Journal of Finance* 67, 431-477.
- Pastor, Lubos, and Pietro Veronesi, 2003, Stock valuation and learning about profitability, *Journal of Finance* 58, 1749-1789.
- Pastor, Lubos, and Pietro Veronesi, 2006, Was there a Nasdaq bubble in the late 1990s? *Journal of Financial Economics* 81, 61-100.
- Russo, J. Edward, and Paul J. H. Schoemaker, 1992, Managing overconfidence, *Sloan Management Review* 33, 7-17.
- Scheinkman, Jose A., and Wei Xiong, 2003, Overconfidence and speculative bubbles, *Journal of Political Economy* 111, 1183-1219.
- Shiller, Robert J., 2000, *Irrational Exuberance* (Princeton University Press, Princeton).
- Timmermann, Allan, 1993, How learning in financial markets generates excess volatility and predictability in stock prices, *Quarterly Journal of Economics* 108, 1135-1145.
- Timmermann, Allan, 1996, Excess volatility and predictability of stock prices in autoregressive dividend models with learning, *Review of Economic Studies* 63, 523-557.
- Zhang, Lu, 2005, The value premium, *Journal of Finance* 60, 67-103.

Table I: Median Firm Characteristics in the Historical Sample

The table reports median firm characteristics in characteristic-sorted portfolios. Market-to-book assets MB is the market value of assets (market value of equity lagged six months plus book value of debt, where the latter is computed as total liabilities plus preferred stocks minus the sum of deferred taxes and convertible debt) divided by total book assets. Asset growth AG is the percentage change in book assets from the last year to this year. Operating profitability OP is operating income before depreciation (total revenue minus cost of goods sold minus selling, general and administrative expenses) divided by total book assets. Portfolios are formed based on quintile values of the sorting characteristic.

	Portfolios Sorted on OP			Portfolios Sorted on AG			Portfolios Sorted on MB		
	OP	AG	MB	OP	AG	MB	OP	AG	MB
L	0.046	0.075	1.351	0.118	-0.036	1.165	0.109	0.048	0.896
2	0.110	0.082	1.103	0.139	0.042	1.185	0.129	0.075	1.112
3	0.145	0.089	1.228	0.154	0.100	1.333	0.149	0.099	1.376
4	0.186	0.109	1.517	0.167	0.186	1.572	0.178	0.131	1.874
H	0.262	0.149	2.323	0.151	0.483	1.806	0.209	0.208	3.424

Table II: Portfolio Returns in the Historical Sample

The table reports mean returns, volatilities, and Sharpe ratios of long-short portfolios. All statistics are computed using monthly data and reported in annualized terms. Mean returns and volatilities are percentages. Market minus risk-free is the market portfolio return in excess of the risk-free rate. Value minus growth, low minus high asset growth, and high minus low profitability are based on the *MB*, *AG*, and *OP* sorts, respectively. All portfolios are value-weighted. Market-neutral returns are computed as the alpha plus the residual in monthly time-series CAPM regressions of the portfolio return on the value-weighted market return. Industry-adjusted portfolio returns are measured in excess of the 48 Fama-French industry portfolio returns.

Long-Short Portfolios	Raw Returns			Market-Neutral Returns			Industry-Adjusted Market-Neutral Returns		
	Mean Return	Volatility	Sharpe Ratio	Mean Return	Volatility	Sharpe Ratio	Mean Return	Volatility	Sharpe Ratio
Market minus Risk-Free	5.88	15.62	0.376	-	-	-	-	-	-
Value minus Growth	4.07	13.74	0.297	5.17	13.42	0.385	3.03	5.76	0.526
Low minus High Asset Growth	3.76	11.66	0.322	5.51	10.69	0.515	3.63	5.96	0.610
High minus Low Profitability	3.71	13.22	0.281	5.40	12.43	0.435	3.84	9.04	0.425

Table III
Model Parameters

The table reports the parameter values in the base-case calibration. More detailed descriptions of the parameters and their calibration are provided in Section II.B.

Parameter	Value
Discount rate r	0.050
Mean reversion rate of disruption climate ρ_μ	0.070
Volatility of disruption climate σ_μ	0.100
Volatility of transitory disruption shocks σ_M	0.250
True signal precision η	0.500
Biased signal precision η_B	0.934
Expected time in the early growth state $1/q_{EG}$	3 years
Expected time in the mature growth state $1/q_{MG}$	4 years
Expected time in the no growth state $1/q_{NG}$	3 years
Project capital investment k	0.950
Project profitability in the early growth state a_{EG}	0.250
Project profitability in the mature growth state a_{MG}	0.125
Average project termination rate λ	0.150
Capital recovery rate α	0.650

Table IV
Summary Statistics of Simulated Data

The table reports summary statistics of simulated data samples in the base-case calibration. See Section III.A for further details of the reported statistics.

Panel A: Full sample							
	Market portfolio return (Annual %)	Median capital stock K	Median profitability f/K	Median Tobin's q V/K			
Mean across all simulated samples	5.00	3.150	0.144	1.371			
Standard deviation of sample means	0.75	0.232	0.005	0.055			
Mean of sample standard deviations	2.61	0.267	0.004	0.077			
Panel B: Median firm characteristics							
	Capital stock K	Profitability f/K	Asset growth (%)	Tobin's q V/K			
Early Growth	2.418	0.171	22.35	2.242			
Mature Growth	4.057	0.138	9.21	1.285			
No Growth	2.670	0.138	-13.91	1.192			
Panel C: Characteristic portfolios							
	% Early Growth	% Mature Growth	% No Growth	Capital stock K	Profitability f/K	Asset growth (%)	Tobin's q V/K
Value	0.0	33.5	66.5	2.515	0.113	-12.22	1.106
Growth	100.0	0.0	0.0	1.911	0.138	37.54	2.312
Low asset growth	0.1	0.1	99.8	1.673	0.122	-13.94	1.299
High asset growth	57.5	39.8	2.7	2.007	0.122	37.60	2.249
Low profitability	38.5	34.8	26.7	1.579	0.092	12.15	1.296
High profitability	58.1	24.0	17.9	3.884	0.212	8.53	2.049

Table IV – continued

Panel D: Characteristic strategy returns

	Volatility (Annualized %)	Correlations		
		Value minus Growth	Low minus High Asset Growth	High minus Low Profitability
Value minus Growth	3.85	1	0.955	0.914
Low minus High Asset Growth	3.49	-	1	0.889
High minus Low Profitability	3.36	-	-	1

Table V
Sharpe Ratio Distributions

The tables reports percentiles of Sharpe ratio distributions. Panel A reports percentile values of the theoretical benchmark distribution described in Section III.B and the corresponding z -values. Panels B and C report percentile values of the Sharpe ratio distributions of various portfolios in simulations with rational and biased investors, respectively.

	Percentile						
	1 st	5 th	10 th	50 th	90 th	95 th	99 th
Panel A: Benchmark							
z -value	-2.326	-1.645	-1.282	0	1.282	1.645	2.326
Sharpe Ratio	-0.329	-0.233	-0.181	0	0.181	0.233	0.329
Panel B: $\eta_B = \eta = 0.500$							
Disruption Surprise Factor	-0.328	-0.231	-0.179	0.002	0.181	0.231	0.333
Value minus Growth	-0.361	-0.265	-0.212	-0.034	0.155	0.211	0.311
Low minus High Asset Growth	-0.383	-0.285	-0.231	-0.056	0.120	0.171	0.268
High minus Low Profitability	-0.329	-0.243	-0.196	-0.032	0.126	0.169	0.237
Panel C: $\eta_B = 0.934$							
Disruption Surprise Factor	-0.750	-0.535	-0.421	0.002	0.424	0.535	0.758
Value minus Growth	-0.712	-0.505	-0.402	-0.012	0.367	0.470	0.656
Low minus High Asset Growth	-0.743	-0.533	-0.428	-0.039	0.326	0.428	0.610
High minus Low Profitability	-0.568	-0.405	-0.322	0.008	0.332	0.422	0.571

Table VI
Serial Correlations and Variance Ratios of Returns

The table reports serial correlations and variance ratios of returns in simulations with biased investors. The reported mean values, and the standard deviations in parentheses, are computed across the 10,000 simulated samples. Panel A reports serial correlations of returns measured at intervals that range from one month to five years. Fractions of simulated samples with positive serial correlations are reported in brackets. Panel B reports variance ratios of returns, where the numerator is the annualized variance of returns measured at intervals that range from one year to five years and the denominator is the annualized variance of monthly returns. Fractions of simulated samples where the variance ratio exceeds one are reported in brackets.

	Return horizon	Disruption Surprise Factor	Value minus Growth	Low minus High Asset Growth	High minus Low Profitability
Panel A: Serial correlation	One month	0.024 (0.044) [0.712]	0.020 (0.050) [0.654]	0.019 (0.053) [0.645]	0.014 (0.061) [0.594]
	One year	0.167 (0.148) [0.867]	0.129 (0.152) [0.799]	0.127 (0.155) [0.791]	0.076 (0.167) [0.663]
	Two years	0.206 (0.202) [0.838]	0.165 (0.203) [0.783]	0.160 (0.209) [0.775]	0.098 (0.211) [0.669]
	Five years	0.103 (0.295) [0.641]	0.073 (0.291) [0.606]	0.069 (0.293) [0.601]	0.022 (0.288) [0.530]

Table VI - continued

	Return horizon	Disruption Surprise Factor	Value minus Growth	Low minus High Asset Growth	High minus Low Profitability
Panel B: Variance ratio	One year	1.246	1.198	1.190	1.129
		(0.240)	(0.249)	(0.247)	(0.265)
		[0.856]	[0.786]	[0.770]	[0.660]
	Two years	1.475	1.372	1.357	1.233
		(0.440)	(0.436)	(0.423)	(0.442)
		[0.883]	[0.813]	[0.802]	[0.675]
	Five years	1.958	1.737	1.705	1.445
		(0.969)	(0.895)	(0.865)	(0.831)
		[0.871]	[0.808]	[0.796]	[0.664]

Table VII
Sharpe Ratios of Implementable Investment Strategies

The table reports percentile values of the distributions of Sharpe ratios generated by implementable investment strategies described in Section III.C.

	Sharpe Ratio Distribution				
	Percentile				
	5 th	10 th	50 th	90 th	95 th
Panel A: Factor Timing Strategies					
Optimal Factor Timing Strategy	0.073	0.146	0.413	0.703	0.789
Quantitative Factor Timing Strategy	0.067	0.138	0.407	0.696	0.781
Panel B: Characteristic Timing Strategies					
10-year Estimation Period:					
Value	-0.113	-0.047	0.216	0.507	0.593
Asset Growth	-0.110	-0.044	0.219	0.510	0.597
Profitability	-0.140	-0.078	0.159	0.414	0.488
25-year Estimation Period:					
Value	-0.211	-0.142	0.122	0.425	0.518
Asset Growth	-0.212	-0.144	0.123	0.423	0.517
Profitability	-0.211	-0.148	0.084	0.344	0.422
50-year Estimation Period:					
Value	-0.294	-0.215	0.058	0.389	0.492
Asset Growth	-0.294	-0.218	0.056	0.387	0.489
Profitability	-0.275	-0.207	0.034	0.316	0.398

Table VIII
Comparative Statics and Alternative Specifications

The table reports the results from model comparative statics and alternative specifications that are described in Section III.D.

Sharpe Ratio Distribution Percentile	Value minus Growth Portfolio		Optimal Factor Timing Strategy		Value Timing Strategy – 10-year Estimation Period	
	90 th	95 th	50 th	90 th	50 th	90 th
Panel A: Base Case	0.367	0.470	0.413	0.703	0.216	0.507
Panel B: Overconfidence						
Low, $\eta_B = 0.869$	0.273	0.354	0.286	0.528	0.119	0.359
High, $\eta_B = 0.981$	0.550	0.707	0.594	0.976	0.379	0.765
Maximum, $\eta_B = 1$	0.790	1.016	0.774	1.283	0.549	1.059
Panel C: Persistence of Disruption Shocks						
One-year half-life, $\rho_\mu = 0.693$	0.171	0.228	0.149	0.333	-0.018	0.169
Five-year half-life, $\rho_\mu = 0.139$	0.301	0.386	0.352	0.611	0.145	0.398
15-year half-life, $\rho_\mu = 0.046$	0.380	0.494	0.430	0.727	0.234	0.536
Panel D: Initial Optimism						
Unbiased signal precision	0.240	0.295	0.337	0.479	0.023	0.212
Biased signal precision	0.578	0.684	0.634	0.914	0.313	0.625

Figure 1
Sharpe Ratio Distributions of Characteristic-Sorted Portfolios with Unbiased Beliefs




Figure 2
Sharpe Ratio Distributions of the Value minus Growth Portfolio

