Innovation, Creative Destruction, and the Cross-Section of Stock Returns

AYDOĞAN ALTI and SHERIDAN TITMAN[1]

Since the late 1970s, financial economists have been aware of what has been referred to as asset pricing anomalies.[2] In brief, the literature has documented that various past return measures, price scaled (value) variables (e.g., the price to book assets ratio), and profitability measures (e.g., gross profits to assets) have historically predicted future returns.[3] Although past return strategies have historically generated high Sharpe ratios (e.g., short-term return reversals and momentum) the focus of this study is on anomalies that are based on firm fundamentals, which are consistent with greater mispricing and are slower to correct.

As shown in Table 1, an industry neutral value-weighted portfolio that optimally tilts towards high book to market and highly profitable stocks achieves a Sharpe ratio nearly three times the Sharpe ratio of the market portfolio over the past 50 years, (see also Novy-Marx (2013) for similar findings using portfolios that are not constructed to be industry neutral). These historical return patterns are consistent with three possibilities: The first is that the returns reflect systematic risks that investors wish to avoid. While there is a large literature that considers this possibility, the magnitude of the historical Sharpe ratio of the combined value and profitability

[1] University of Texas at Austin.
[2] See the 1978 special issue of the *Journal of Financial Economics* (Vol. 6, Issues 2-3) on anomalous evidence regarding market efficiency.
[3] Add specific cites in this footnote.

strategy appears to be too large to be consistent with plausible preferences.[4] The second possibility is that history has provided an unusual sample path simply by chance.[5] The very high t-statistics associated with the returns we observe suggest that this explanation is also unlikely. Finally, the return patterns may reflect systematic mistakes that investors may have made. This third possibility is the focus of our study.

To explore the types of mistakes investors can make we develop a model where investors learn about a latent state variable that influences a source of systematic risk in the economy. The systematic source of risk that we examine is what we will refer to as the innovation climate, which influences the arrival rate of new investment projects. Because potential mistakes are made about a systematic source of risk, there is no arbitrage in our proposed economy.[6] Indeed, the analysis is consistent with the cross-sectional implications of multi-factor models, like Fama and French's (1993, 2015) three and five factor models. However, the return premia associated with the risk factors in these models can potentially be either too high or too low relative to what we would observe in a setting with fully rational investors.

To understand our model, and how it connects with the empirical evidence, it is useful to think about firms as combinations of assets in place and growth opportunities. Firms in this economy differ along two dimensions. The first dimension describes their access to new growth opportunities; growth firms are endowed with new projects every period while value firms

[4] There is a large literature that discusses the magnitude of the Sharpe ratio of the overall market, but we are not aware of existing research that tries to rationalize the much larger magnitude of the Sharpe ratio of combined value/profitability strategies. 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.

[5] A related explanation is that the return anomalies were discovered with the benefits of data mining. While this is always a possibility, the value anomaly appears to be very robust. It is natural to ask whether stocks with high values relative to measures of fundamental value are over or under-valued and the excess returns of value scaled variables holds for all the plausible measures of fundamental value. The profitability anomaly is not nearly as robust, and in this respect, the data-mining critique is more applicable. Indeed, as we will discuss later in the paper, although Sharpe ratios of gross profitability sorted portfolios are quite high, portfolios sorted on net profitability do not generate particularly high Sharpe ratios.

[6] In contrast, there can be arbitrage opportunities for rational investors when most investors learn slowly about firm specific information.

simply harvest the profits from their existing opportunities. The second dimension is the firm's history. New growth firms, like Twitter, have very little in the way of assets in place and thus low profits, while more mature growth firms, like Apple, have very profitable assets in place as well as growth opportunities.

Our model abstracts from an overall productivity (i.e., the market) factor and focuses on the innovation climate, which can have a technical component (e.g., electrification in the 1920s and the internet in the 1990s) as well as regulatory and policy components that influence the arrival of new investment opportunities. A favorable climate increases the arrival rate of new projects, benefiting the young growth firms and some of the mature growth firms. However, since these new projects compete with existing businesses, a favorable climate is associated with declines in 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.

The innovation shocks can have temporary as well as persistent components and investors imperfectly distinguish between these components from the observed arrival rate of new projects and from soft information about technological advances and the political climate. A positive shock to innovation benefits growth firms with very few existing projects, since it increases the rate at which it realizes new opportunities, and hurts value firms, since the new projects will compete with their existing projects. The market will over-react (relative to the case with full information) to this shock if it turns out that a larger than anticipated portion of the shock is temporary, and will under-react if it turns out that the shock is mostly permanent.

Since the innovation prospects in the economy are unobservable and change from year to year even fully rational investors are sometimes too optimistic and sometimes too pessimistic about the permanence of an innovation shock. For example, in the late 1990s, investors may have been too optimistic about the long-term growth prospects generated by the internet, making them too pessimistic about the long-term viability of profitable old economy firms. When this is the case, growth stocks are overvalued and more profitable stocks are undervalued. Hence, within the context of our model, value stocks outperform growth stocks and the stocks of more profitable firms outperform their less profitable counterparts along sample paths with less than anticipated innovation.

The model is flexible enough to consider full rationality, biased pre-conceptions, and slow learning that arise, for example, from over-confidence. As we show, even under complete rationality, one can generate value and profitability effects in small samples. However, if investor priors about the unknown parameter are drawn from a distribution that is centered on the true distribution generating the parameter, and if these priors are updated using Bayes Theorem, then sample paths that reject market efficiency at the 5% level occur about 5% of the time. In other words, parameter uncertainty does not by itself lead to biased inferences, implying that one needs more than just parameter uncertainty to explain these asset pricing anomalies; some degree of irrationality is needed. Moreover, because investors do learn, the model does not generate value and/or profitability anomalies in sufficiently long time series even if investors are not fully rational. However, depending on their initial priors and how quickly investors learn, these anomalies can be generated in small samples. A contribution of our quantitative model is that allows us to more precisely define and explore what we mean by long versus short samples and slow versus rational learning.

We consider behavioral biases that arise from two sources. The first source is the tendency of investors to be overly optimistic about the impact of new technology. This bias, which can influence the initial beliefs of investors, can generate the value and profitability anomaly in the early part of our sample. However, because investors do learn in our model, this effect should diminish with time. The second source is a tendency of investors to be overconfident about their abilities to evaluate soft information. This tendency does not necessarily bias investors towards any particular type of stock; however, because it does slow down the learning process, it increases the probability of observing pricing anomalies in small samples. For example, sample paths that are only expected to occur 5% of the time under full rationality may occur 30% of the time when investors learn slowly.

The analysis in this paper is closely related to a growing literature that examines asset pricing in settings where rational investors learn about uncertain parameters. For example, Lewellen and Shanken (2002) show, within a setting where investors learn about unknown parameters, that returns may look predictable ex post, even when they are not at all predictable ex ante. Similarly, Pastor and Veronesi (2003, 2006) explore settings with rational learning where return patterns resemble bubbles that inflate and subsequently deflate. What these papers show is that there exist sample paths where the choices of rational investors generate unusual return patterns. However, as we mentioned above, our analysis shows that these sample paths are not generated "too often" when investors are rational. Even with parameter uncertainty, some form of irrationality is required for the null to be rejected more than would be expected by pure chance.

Our analysis is also related to the behavioral finance literature. For example, Daniel, Hirshleifer, and Subrahmanyam (1999) described a link between the value effect and the tendency of investors to be overconfident about the precision of their private information, and Barbaris, Shleifer and Vishny (1998) considered behavioral biases that influence how investors interpret information about the persistence of earnings shocks. We contribute to this literature by providing a model of mispricing of systematic sources of risk and by explicitly incorporating dynamic learning. Moreover, our quantitative model can be used to simulate returns that can be directly compared to the actual time-series pattern of stock returns.

Our model also contributes to a series of papers that explore how investor perceptions can endogenously generate covariation between stocks. For example, Campbell and Vuolteenaho (2004, AER) explore the observation that part of the covariation amongst growth stocks, and an important component of their market beta, is due to fluctuations in their discount rates. In our model, overconfident investors tend to over-react to soft signals that pertain to the permanence of innovation shocks, and in doing so they induce excess (relative to the full information case) covariation amongst value and growth stocks. In this sense, our model endogenously generates what looks like a sentiment factor that induces excess covariation.

Our model also complements models by Kogan, Papanikolaou, and Stoffman (2015) and others, which also examines how the innovation process can generate sources of systematic risk that affect the prospects of different firms differently. Given the magnitude of the observed Sharpe ratios, these models would require fairly extreme risk preferences to explain the observed return patterns over the past 50 years. Although we solve our model with risk neutral preferences, effectively shutting down the risk aversion channel, we can envision a model that accounts for risk preferences as well as slow learning that better explain the observed return patterns.