Empirical Cross-Sectional Asset Pricing*

Stefan Nagel
University of Michigan, NBER and CEPR

July 9, 2013

Abstract

I review recent research efforts in the area of empirical cross-sectional asset pricing. I start by summarizing the evidence on cross-sectional return predictability and the failure of standard (consumption) CAPM models and their conditional versions to explain these predictability patterns. One response in part of the recent literature is to focus on ad-hoc factor models, which summarize the cross-section of expected returns in parsimonious form, or on production-based approaches, which suggest links between firm characteristics and expected returns. Without imposing restrictions on investor preferences and beliefs, neither one of these two approaches can answer the question why investors price assets the way they do. Within the rational expectations paradigm, recent research that imposes such restrictions has focused on the ICAPM, long-run risks models, as well as frictions and liquidity risk. Approaches based on investor sentiment have focused on the development of empirical proxies for sentiment and for the limits to arbitrage that allow sentiment to affect prices. Empirical work that considers learning and adaptation of investors has worked with out-of-sample tests of cross-sectional predictability.

*I am grateful for comments to Frederico Belo, John Campbell, Marco Giacoletti, Sebastian Infante, Bob Hodrick, Arthur Korteweg, Ian Martin, Dimitris Papanikolaou, Lasse Pedersen, Jeff Pontiff, Mike Schwert, and Jeff Wurgler.

1 Introduction

It is a fundamental question of financial economics whether assets with different risk exposures and characteristics are priced to earn different rates of return. I survey studies in empirical cross-sectional asset pricing that provide evidence on this question and relate the evidence to models of investor preferences and beliefs. I begin by laying out a framework that will be helpful for categorizing and interpreting empirical work in this area.

Consider an economy whose state at time t is summarized by a state vector X_t. A vector of parameters ρ pins down the law of motion of X_t.[1] Assume that X_t is observable to investors at t, and let $\mathcal{J}_t$ denote the information set generated by investors' observations of X. The economy features N assets that produce a vector of payoffs $D_t(X_t)$. Let $M_t(X_t, \phi)$ be the stochastic discount factor (SDF) that reflects preferences of a representative investor with preference parameters ϕ. This investor is endowed with knowledge of ϕ, but not necessarily of ρ. The investor's degree of knowledge of ρ is important for asset pricing, because it affects her ability to forecast future realizations of X, and hence future asset payoffs and the SDF, which depend on X.

Most empirical studies in cross-sectional asset pricing rely on rational expectations asset-pricing theory in the tradition of Lucas (1978) to derive model predictions and a null hypothesis. Under rational expectations (Muth (1961)), investors are endowed with knowledge of the parameters ρ. The econometric framework in much of empirical asset pricing is built on this strong assumption. Defining (gross) returns as $R_{t+1} \equiv (P_{t+1} + D_{t+1})/P_t$, investors price assets according to

$$E[R_{t+1}(X_{t+1})M_{t+1}(X_{t+1}, \phi)|\mathcal{J}_t] = 1, \tag{1}$$

where the conditional expectation $E[.|\mathcal{J}_t]$ is taken under the density $f(X_{t+1}|\mathcal{J}_t, \rho)$ with knowledge of ρ. Now define $H_{t+1} \equiv R_{t+1}M_{t+1} - 1$. Continuing to assume knowledge of ρ, we can apply the law of iterated expectations and condition down to the econometrician's information

[1] For example, if X follows a univariate AR(1) with normally distributed innovations, ρ includes the mean of X, the autoregressive coefficient, and the variance of innovations.

set, $\mathcal{A}_t$, where $\mathcal{A}_t \subseteq \mathcal{J}_t$. The resulting conditional moment restrictions imply orthogonality conditions

$$E[H_{t+1}(X_{t+1}, \phi) \otimes Z_t] = 0 \tag{2}$$

for any $\mathcal{A}_t$-measurable instruments Z_t. Studies in cross-sectional asset pricing often focus on the special case

$$E[H_{t+1}(X_{t+1}, \phi)' Z_t] = 0 \tag{3}$$

of (2), where Z_t is a vector of portfolio weights that combines the N assets into a portfolio. Evaluating (3) is akin to asking whether the elements Z_{it} of Z_t can predict differences in discounted returns H_{it+1} of the N assets. The cross-sectional asset pricing literature has produced a large body of evidence of such cross-sectional return predictability. I start the survey with a review of this evidence in Section 2. The economic interpretation of these findings centers around the following key issues:

Stochastic discount factor. The law of one price implies that a stochastic discount factor (SDF) exists such that (1) holds, and it can be constructed as a linear combination of the asset payoffs (Hansen and Richard (1987)). The economic content of pricing models is therefore in the restrictions that they impose on the SDF. Restrictions can be motivated by theory (e.g., the CAPM) or have an ad-hoc nature (e.g., restricting a factor model to a small number of ad-hoc factors). Researchers often look to enrich the specification of investors' preferences or the dynamics of risks to enhance the pricing performance of models. I review these issues in Sections 3 and 4.

Learning. The standard approach to estimating and evaluating asset-pricing models is based on the rational expectations paradigm, where it is presumed that investors know the parameters ρ of the law of motion of X_t. This is an extremely demanding notion. More realistically, an investor would have to learn ρ by observing data. A Bayesian investor learning ρ would form expectations based on the predictive density $f(X_{t+1}|\mathcal{J}_t) = \int f(X_{t+1}, \rho|\mathcal{J}_t) d\rho$, viewing ρ as a random variable, and pricing assets according to

$$E^B[R_{t+1}(X_{t+1})M_{t+1}(X_{t+1},\phi)|\mathcal{J}_t] = 1, \tag{4}$$

where the B superscript of the expectations operator indicates that the expectation is now taken under the Bayesian predictive density $f(X_{t+1}|\mathcal{J}_t)$. For the econometrician, this has the unpleasant consequence that the usual approach of evaluating a sample version of (2) to test the pricing model is not valid: return predictability that is evident to an econometrician ex-post may not exist under the predictive distribution that investors perceive at the time they make decisions. There is only a small empirical literature so far that looks at empirical cross-sectional asset pricing from the perspective of learning. Out-of-sample tests are one—albeit imperfect—way to study learning effects, and I cover this line of work in Section 6.

Sentiment. The probability distribution of X_{t+1} perceived by investors at time t may deviate from the rational-expectations and Bayesian-learning benchmarks. In this case

$$E^S[R_{t+1}(X_{t+1})M_{t+1}(X_{t+1},\phi)|\mathcal{J}_t] = 1 \tag{5}$$

holds with $E^S[.|\mathcal{J}_t]$ evaluated under the subjective density perceived by investors. I use the label "sentiment" as a catch-all term to describe expectations-formation of this kind. Section 5 covers empirical studies that take this perspective. Since the boundary between rational theories and sentiment-approaches is fuzzy,[2] I refrain from using the somewhat simplistic categorization of cross-sectional asset-pricing studies into "rational" and "irrational", or "behavioral," approaches.

Frictions. Frictions can affect the SDF in several ways. Frictions may prevent participation of some classes of investors in the market. Frictions can also give rise to risks that are of concern to investors, such as liquidity risk. I cover work on these issues in parts of Section

[2]For example, macroeconomists have worked with models of boundedly rational learning (Bray (1982), Marcet and Sargent (1989)) in which agents use simplified forecasting rules rather than a full Bayesian approach. Sims (2003) argues that agents with limited information-processing capacity may find it rational to be inattentive to some types of information.

4. Frictions also play an important role in the sentiment-based approaches in Section 5 as an explanation as to why sophisticated investors—whose beliefs may be closer to the Bayesian-learning or rational-expectations benchmarks—may be unable to exploit and eliminate the effects of sentiment on asset prices.

Data snooping. One problem that plagues cross-sectional asset pricing is that much of the empirical work in this area is inherently a search for anomalies. When researchers search over a large number of candidate predictors for predictability that is unexplained by standard models of the SDF, it is inevitable that some turn up as "significant" just by pure chance. Conventional procedures of statistical inference do not account for this search over many candidate predictors. The out-of-sample tests reviewed in Section 6 can help shed light on this problem.

Market efficiency. Empirical work on cross-sectional return predictability is often discussed as the study of market efficiency, but the meaning of this concept is not clearly defined, which limits its usefulness as an organizing principle. Fama (1970) defines a (semi-strongly) efficient market as one that always "fully reflects" all available (public) information, but it is not clear whether "fully reflect" refers to the demanding notion of rational expectations (this is the view taken in Jensen (1978)), or whether the notion of market efficiency should allow for learning and adaptation (Schwert (2003) and Timmermann and Granger (2004) propose this interpretation). Frictions and risks further obscure the precise meaning of market efficiency. Jensen (1978) augments his rational-expectations definition of market efficiency by suggesting that market efficiency should be defined as the "absence of profit opportunities" from trading on public information. Yet, in a market in which sophisticated and sentiment-driven investors trade, it is possible that sentiment affects prices, yet sophisticated investors do not perceive any abnormal profit opportunities given the frictions and risks that they face.[3] If one defines such a market as efficient, efficiency becomes a largely empty concept because it does not distinguish between interesting alternative hypotheses (rational expectations, learning,

[3] See, for example, Dumas, Kurshev, and Uppal (2009) for an illustration of this point in a general equilibrium model.

sentiment) of how investors price assets. For these reasons, I prefer to explicitly spell out the belief-formation hypothesis that researchers work with, rather than relying on the imprecise notion of market efficiency.

While I aim to cover the main conceptual issues in empirical cross-sectional asset pricing, the survey is necessarily selective in its scope. I focus on recent work, mostly on equity markets. On the econometric side, I emphasize the importance of the choice of moment conditions when estimating and evaluating asset-pricing models, but I mostly skip a discussion of methods for estimation and testing.[4]

2 Cross-Sectional Return Predictability

In this section, I provide a brief survey of the vast literature on cross-sectional return predictability. I focus on the most common categories of return predictors. Most studies in this literature assess predictability relative to standard linear factor models of the SDF, such as the CAPM, consumption CAPM, or ad-hoc factor models as benchmark models. I discuss these models later in Section 3.

2.1 Technical Predictors

The history of returns and trading volume of a stock is a natural place to look for return predictors, and the data are easily accessible. In the practice of investment management, predictors of this kind are in the realm of "technical analysis."

One of the most studied predictability patterns in this area is the momentum effect that entered the academic literature with Jegadeesh and Titman (1993). Jegadeesh and Titman show that stocks with high returns over the past three to twelve months (winners) outperform stocks with low recent returns (losers) over the next three to twelve months. A related phenomenon is the post-earnings-announcement drift. Stocks with unexpectedly

[4]Recent surveys that devote more attention to the methodology for estimation and testing include Jagannathan, Skoulakis, and Wang (2009) and Goyal (2012).

good earnings outperform those with unexpectedly bad earnings over the next six months (Bernard and Thomas (1989)). Chan, Jegadeesh, and Lakonishok (1996) show that this earnings momentum is related to, but partly distinct from, the price momentum effect. Lee and Swaminathan (2000) show that stocks with high trading volume have low future returns and stronger momentum effects.

Over longer horizons, in contrast, returns have a tendency to mean-revert, as shown in DeBondt and Thaler (1985) and DeBondt and Thaler (1987): Winners over the past three to five years underperform losers. Jegadeesh and Titman (2001) find that the momentum effect dies out after a holding period of about 12 months, and returns then start to revert at longer holding periods.

2.2 Valuation Ratios and Profitability

A second category of predictors can be motivated within a simple present-value framework. Vuolteenaho (2000) and Cohen, Polk, and Vuolteenaho (2003) build on the log-linear present value model of Campbell and Shiller (1988) to formulate a present-value relationship in terms of the book-to-market (B/M) ratio. They start with the clean-surplus accounting relation $D_t = Y_t - \Delta B_t$, where D_t is dividends, Y_t is earnings, and ΔB_t is the change in book value of equity. Assuming no equity issues or repurchases and a covariance-stationary B/M, they obtain a log B/M ratio

$$bm_t \approx E_t \sum_{j=0}^{\infty} \rho^j (r_{t+j} - y_{t+j}), \tag{6}$$

where r_t denotes the log stock return, and $y_t \equiv \log(1 + Y_t/B_{t-1})$ is the log return on equity (ROE), and ρ is a constant close to one.[5]

Thus, holding fixed expectations of future ROE, expected returns must be low if bm is low. Holding fixed current bm, high expected ROE implies high expected returns. This motivates the use of bm and expected profitability, preferably jointly, as return predictors.

[5] If the firm also engages in equity issues or repurchases, y_t must be modified to include growth in per-share book value of equity due to the issuance or repurchase activity (Ohlson (2005)). Equity issues at $B/M < 1$ or repurchases at $B/M > 1$ both raise the per-share book value of equity.

Of course, expected returns do not necessarily have to vary across stocks. If they do not, variation across stocks in bm reflects either future profitability or the approximation error in (6). But if they do vary, bm and/or expected ROE should be correlated with expected returns.

As (6) is simply an approximate accounting identity, it does not shed light on the structural drivers of cross-sectional variation in expected returns. It does not matter whether priced risk under rational expectations, learning effects, or sentiment drive the variation in expected returns. The conclusion that bm and expected profitability should predict returns applies in each one of these cases.

The use of valuation ratios like B/M in return prediction has a long history in finance. Fama and French (1992) reexamine the evidence and confirm that stocks with low market capitalization (small firm premium) and high B/M (value premium) earn higher expected returns. They find that firm size and B/M absorb the predictive role of the earnings-to-price ratio. Other researchers examine B/M jointly with current profitability measures as proxies for expected future profitability. Haugen and Baker (1996) document that ROE positively predicts returns controlling for B/M and various other return predictors. Vuolteenaho (2002) shows that recent returns, ROE, and B/M jointly predict returns. Given the approximate accounting identity, it is not all that surprising that valuation ratios like B/M and profitability variables are positively related to future returns. The more interesting bit is the failure of standard models of risk and return to explain this cross-sectional variation in returns.

Current ROE may not be the best predictor of future ROE. For example, some subcomponents of current profitability may help predict future profitability better than others. Sloan (1996) decomposes earnings into cash earnings and accruals. He finds that firms with a high accruals have low future profits and abnormally low future returns. Hirshleifer, Hou, Teoh, and Zhang (2004) argue that net operating assets summarize the cumulative history of non-cash earnings and provide a better predictor of future earnings than single-period accruals. They find that net operating assets strongly predict future returns. Fama and French

(2006) employ a variety of accounting variables, past stock returns, and analyst earnings forecasts that help forecast future earnings and they show that these variables also predict returns. Novy-Marx (2012a) argues that gross profits (revenues minus cost of goods sold) are a good predictor of future earnings, and he finds that gross profitability is strongly positively related to future stock returns.

Some researchers have used valuation ratios, profitability measures, and other accounting variables to construct a summary measure of financial distress and find that distressed stocks have abnormally low returns (Dichev (1998) and Griffin and Lemmon (2002)). Campbell, Hilscher, and Szilagyi (2008) report a similar finding within a distress-prediction framework that also employs price-based variables such as lagged returns and volatility.

2.3 Firm Investment and Financing

Theories of firms' investment and financing can provide further insights about potential predictors of returns. According to the q-theory of investment, firm investment is positively related to expected profitability and negatively related to discount rates and hence future stock returns. Li, Livdan, and Zhang (2009) provide a recent exposition of this mechanism. Titman, Wei, and Xie (2004), Fama and French (2006), Anderson and Garcia-Feijóo (2006), and Li, Livdan, and Zhang (2009) find evidence consistent with this prediction.

High-investment firms are also likely to raise external financing, while firms with low investment rates distribute capital to investors. As a result, discount rates are also negatively related to external financing (see, e.g., Li, Livdan, and Zhang (2009)). Consistent with this prediction, Daniel and Titman (2006), Fama and French (2008), and Pontiff and Woodgate (2008) construct measures of net equity issuance activity and find a negative relationship to future stock returns.

High levels of investment and external financing also translate into higher growth. Therefore, one would expect a negative association between growth and future returns. Lakonishok, Shleifer, and Vishny (1994) document a negative relation between past sales growth and returns. Cooper, Gulen, and Schill (2008) look at growth in total assets and find a negative correlation with future returns.

Investment and growth are also related to accruals. Fairfield, Whisenant, and Yohn (2003) argue that high-growth firms tend to be high-accrual firms, suggesting that the relationship between accruals and expected returns of Sloan (1996) may be a growth effect. Lewellen and Resutek (2012) however show that investment, external financing, and accruals all have distinct roles in predicting returns.

While q-theory is useful for identifying potential cross-sectional predictors of returns, the theory is silent about the reasons why investors price stocks to have different expected rates of return. The evidence that investment- and financing-related variables predict returns does not reveal whether priced risk under rational expectations, investor learning, or sentiment are the drivers of the discount rates that firms face in financial markets and that they respond to in their investment and financing decisions.

2.4 Idiosyncratic risk

Theories in which idiosyncratic volatility plays a role in pricing typically predict a positive relation between idiosyncratic volatility and expected returns. For example, in Merton (1987) a positive relationship arises because investors are imperfectly diversified and demand compensation for bearing idiosyncratic volatility. Empirically, Ang, Hodrick, Xing, and Zhang (2006) find the opposite result: stocks with high idiosyncratic volatility have extremely low returns. Huang, Liu, Rhee, and Zhang (2010) caution that this idiosyncratic volatility effect seems to be driven by the return reversals known to exist in one-month returns (Jegadeesh (1990)). Idiosyncratic volatility is also one of the variables used in the distress-prediction model of Campbell, Hilscher, and Szilagyi (2008) and the sign of the effect in their framework is consistent with the findings of Ang et al.

2.5 Seasonality

One puzzling feature of cross-sectional return predicability is that it is subject to strong seasonality. Novy-Marx (2012b) offers an updated view based on recent data. He documents that many of the cross-sectional return predictors reviewed in this section have a strong January seasonal: size and value effects, for example, are concentrated in January, while momentum and profitability-related predictability is weaker in January. The research reviewed in this survey is silent about the source of this seasonality.

3 Linear Factor Models

Most empirical studies in cross-sectional asset pricing employ (log-)linear factor representations of the SDF. In this section, I review the canonical models that appear in most empirical work, while Section 4 looks at recent efforts to improve these models. The work reviewed here and in Section 4 takes a rational expectations approach. To economize on notation, I use $E_t[.]$ as shorthand for $E[.|\mathcal{A}_t]$.

3.1 Canonical Models

Typically, the SDF specifications in cross-sectional asset pricing studies are special cases of the SDF that arises in a representative-agent framework in which the representative investor has the preferences proposed by Epstein and Zin (1989), Epstein and Zin (1991), and Weil (1990). In this model, the log SDF follows

$$m_{t+1} = \theta \log \delta - \frac{\theta}{\psi}\Delta c_{t+1} + (\theta - 1) r_{w,t+1}, \tag{7}$$

where $r_{w,t+1} \equiv \log(W_{t+1}/(W_t - C_t))$ is the log return on wealth, $\Delta c_{t+1} \equiv \log(C_{t+1}/C_t)$ is consumption growth, δ is a time-discount factor, $\theta \equiv (1-\gamma)/[1-(1/\psi)]$, ψ is the elasticity of intertemporal substitution, and γ is a relative risk aversion parameter. With power utility,

$\gamma = 1/\psi$ and $\theta = 1$, and we get the consumption CAPM

$$m_{t+1} = \log \delta - \gamma \Delta c_{t+1}. \tag{8}$$

As Giovannini and Weil (1989) show, if the state vector is identically and independently distributed (IID), we get the CAPM,

$$m_{t+1} = -\gamma r_{w,t+1}. \tag{9}$$

Conditional CAPMs. In (8) and (9), the price of risk, controlled by the parameter γ, is constant over time, which is why these models are often referred to as an *unconditional* CAPMs. Theories of time-varying risk aversion motivate a generalization to time-varying coefficients,

$$m_{t+1} = \phi^0(x_t) - \phi^c(x_t)\Delta c_{t+1}. \tag{10}$$

This is a conditional consumption CAPM. For example, in the external habit model of Campbell and Cochrane (1999), x_t is the surplus-consumption ratio and ϕ_t^0 and ϕ_t^c are nonlinear functions of x_t. A conditional CAPM version of (9),

$$m_{t+1} = \phi^0(x_t) - \phi^w(x_t) r_{w,t+1}, \tag{11}$$

is also frequently applied in empirical work, but the theoretical motivation is less clear in this case. One possibility is that investors myopically maximize, period by period, their expected utility of next-period wealth. In this case, $\phi^0(x_t) = 0$ and $\phi^w(x_t) = \gamma(x_t)$, where $\gamma(x_t)$ is a time-varying relative risk aversion coefficient.

Beta representation. Let r_{it+1} denote the log return on asset i, and σ_{it}^2 its conditional variance. If we combine a log-linear factor model $m_{t+1} = \phi_t^0 - \phi_t^{f\prime} f_{t+1}$ with the assumption that r_{it+1} and f_{t+1} are conditionally jointly normal, evaluation of the log of the conditional pricing restriction (1) yields a linear beta representation,

$$E_t[r_{it+1}] + \frac{1}{2}\sigma_{it}^2 = \mu_t^0 + \beta_{it}'\lambda_t, \tag{12}$$

where $\mu_t^0 \equiv -\phi_t^0 + \phi_t^{f\prime}\mu_t^f - \frac{1}{2}\phi_t^{f\prime}\Omega_t\phi_t^f$, $\beta_{it} \equiv \Omega_t^{-1}\mathrm{Cov}_t(r_{it+1}, f_{t+1})$, $\mu_t^f \equiv E_t[f_{t+1}]$, Ω_t is the conditional covariance matrix of f_{t+1} with a vector ω_t^2 of factor variances on the diagonal, and $\lambda_t \equiv \Omega_t\phi_t^f$ are the factor risk premia. If a conditionally risk-free asset exists with time $t+1$ log return r_t^f, and f_t represents log returns on traded assets, we can apply the conditional pricing relation to the risk-free asset and the factors themselves and solve for ϕ_t^0 and ϕ_t^f:

$$\phi_t^0 = -r_t^f + \phi_t^{f\prime}\mu_t^f - \frac{1}{2}\phi_t^{f\prime}\Omega_t\phi_t^f, \tag{13}$$

$$\phi_t^f = \Omega_t^{-1}(\mu_t^f - \iota r_t^f + \frac{1}{2}\omega_t^2), \tag{14}$$

where ι is a vector of ones. Substituting these back into (12), we obtain

$$E_t[r_{it+1}] + \frac{1}{2}\sigma_{it}^2 = r_t^f + \beta_{it}'(\mu_t^f - \iota r_t^f + \frac{1}{2}\omega_t^2). \tag{15}$$

The relations (13) and (14) highlight that ϕ_t^0 and ϕ_t^f have important implications for the time-series dynamics of returns: they control the conditional risk-free rate and conditional risk premia. While estimation of linear factor models in cross-sectional asset pricing studies often focuses on *cross-sectional* pricing implications, it is important to keep in mind that the SDF parameters also have *time-series* implications for conditional moments.

An alternative way of arriving at a linear beta representation is to exponentiate (10) or (11), carry out a (conditional) first-order Taylor approximation, and normalize the SDF by its conditional mean,[6] which yields

$$M_{t+1} = 1 - \tilde{\phi}_t^{f\prime} F_{t+1}. \tag{16}$$

Applying this SDF to the pricing restriction for excess returns, $E_t[M_{t+1}(R_{t+1} - R_t^f)] = 0$, we can solve for

$$E_t[R_{t+1}] = R_t^f + \beta_t(E_t[F_{t+1}] - R_t^f), \tag{17}$$

which is approximately identical to (15), especially for returns measured over short horizons. To provide a standardized treatment, I discuss factor models within the log-linear/log-normal framework, even though some of the original studies use the linearized SDF (16) applied to simple returns.

Mimicking portfolios. The pricing implications of the SDF remain unchanged if we replace non-traded factors by their (conditional) projection on the test asset returns. More precisely, in the log-normal setting the projection is on the log returns of test assets, and on simple returns in case of the linearized SDF (16).

3.2 Ad-hoc Factor Models

The failure of the (consumption) CAPM to explain cross-sectional return predictability has prompted some researchers to resort to ad-hoc factor models. The motivation for this approach is the insight that risk factors in the SDF can be replaced by their mimicking portfolios. The most prominent example is the Fama and French (1993) three-factor model, in which the SDF is specified as a linear function of a market index return, R_M, the difference in returns between portfolios of small and large stocks, SMB, and the difference in returns between portfolios of high and low B/M stocks, HML. Fama and French (1993) show that this model captures most of the cross-sectional variation in average returns of 25 portfolios sorted by

[6] Any constant c, not only $c = 1$, in (16) would represent a valid normalization of the SDF that produces identical pricing implications for conditional expected excess returns. In the case of a misspecified SDF, though, tests of model misspecification are sensitive to the choice of normalization, as pointed out by Kan and Robotti (2008)

size and B/M.

SMB and HML are constructed from portfolios that span the very same expected return spreads along the size and B/M dimensions that the model is trying to explain. Is the model therefore tautological? Not entirely. The economic content in Fama and French (1993) is in the finding that the three factors explain not only cross-sectional, but also most of the time-series variation in the size and B/M portfolios, with time-series R^2 in excess of 90%. This shows that small firms' returns, for example, correlate a lot more strongly with almost-small firms' returns than with large firms' returns. However, given this strong co-movement, it is not surprising that expected returns of small firms are similar to those of almost-small firms and value stocks' expected returns are similar to those of almost-value stocks in the way captured by the factor model; otherwise a near-arbitrage opportunity would arise. Absence of near-arbitrage opportunities arises in any model under minimal restrictions on preferences and beliefs (e.g., as in Ross (1976) and Cochrane and Saá-Requejo (2000)), and these weak restrictions hold in (plausible) models of sentiment and models of learning just as well as in rational expectations models. Therefore, the fact that the Fama-French model explains size and B/M returns does not discriminate between these competing explanations.

Fama and French (1993) suggest that SMB and HML could mimick pervasive priced macroeconomic risk factors. If so, then SMB and HML should explain other features of the cross-section of expected returns, too, not just size and B/M effects. Fama and French (1996) argue that this is the case because the three-factor model also explains the expected return variation related to earnings/price, cash-flow/price, sales growth, and long-term reversals. However, all of these variables are closely related to B/M. The one anomaly in their tests—momentum—that is not, fails to be explained by the three-factor model.

The empirical evidence accumulated since Fama and French (1996) suggests that there are other important sources of cross-sectional variation in expected returns unrelated to the Fama-French factors. Most of the cross-sectional return predictors reviewed in Section 2 have predictive power after adjusting for exposure to the Fama-French factors. To name a few, the idiosyncratic volatility anomaly (Ang, Hodrick, Xing, and Zhang (2006)), the net equity issuance effect (Daniel and Titman (2006)), and the predictability associated with gross-profitability (Novy-Marx (2012a)) are all present after risk-adjusting returns with the Fama-French factors. Lewellen (2011) looks at portfolios sorted on a composite measure of expected returns that combines many known predictors and his portfolios exhibit large abnormal returns after adjusting for exposure to the Fama-French factors.

Judging from this evidence, the Fama-French factors seem to be a convenient way of summarizing the size and value effects, but not more than that. To capture other dimensions of cross-sectional predictability, too, researchers have proposed alternative factors. Carhart (1997) adds a momentum factor to the Fama-French model. Novy-Marx (2012a) proposes a gross-profitability factor. Hou, Xue, and Zhang (2012) construct a four-factor model with a market factor, a size factor, an investment factor, and an ROE factor. The use of profitability- and investment-related factors is motivated by the present-value relationships and q-theoretic models discussed above in Sections 2.2 and 2.3. Of course, one must keep in mind that present-value relationships and q-theory do not answer the question why investors price stocks to have these differences in expected returns.

3.3 Conditional Factor Models

Much of the early empirical work on size and value premia had been carried out with one-factor models,

$$m_{t+1} = \phi_t^0 - \phi^f f_{t+1}, \tag{18}$$

in which the price of risk is constant. This raises the question whether the unexplained size and value premia could simply reflect the fact that researchers ignored time-varying prices of risk that arise in conditional factor models such as (10) or (11).

To implement tests of conditional factor models, we need to specify the relevant state variables x_t and their functional relationship with the SDF coefficients ϕ_t^0 and ϕ_t^f. Jagannathan and Wang (1996), Cochrane (1996), and Lettau and Ludvigson (2001), and a large number of subsequent papers (see the references in Lewellen and Nagel (2006)) (LN) examine models in which the coefficients are linear in the state variables. Applied to our log-linear SDF here, we get

$$m_{t+1} = \phi^0 + \phi^{0x} x_t - \phi^f f_{t+1} - \phi^{fx} x_t f_{t+1}. \tag{19}$$

Thus, the conditional one-factor model has been turned into an unconditional multi-factor model.[7] Although intuitively appealing, this route turned out to be less fruitful than it initially seemed.

To see why, suppose that a conditional one-factor model $m_{t+1} = \phi_t^0 - \phi_t^f f_{t+1}$ is the true model, but a researcher applies a constant-price-of-risk model like (18). How big is the pricing error that results from this misspecification? Lewellen and Nagel (2006) (LN) perform these calculations for a linearized SDF applied to simple returns. Here I show similar calculations within a log-linear/log-normal framework. The conditional factor model implies

$$E_t[r_{it+1}] - r_t^f + \frac{1}{2}\sigma_{it}^2 = \mathrm{Cov}_t(r_{it+1}, f_{t+1})\phi_t^f. \tag{20}$$

Taking unconditional expectations and noting that the expected conditional covariance equals the unconditional covariance if f is conditionally de-meaned, we get

$$E[r_{it+1}] - E[r_t^f] + \frac{1}{2}E[\sigma_{it}^2] = \mathrm{Cov}(r_{it+1}, f_{t+1} - \mu_t^f)E[\phi_t^f] + \mathrm{Cov}\{\mathrm{Cov}_t(r_{it+1}, f_{t+1}), \phi_t^f\} \tag{21}$$

The first term on the RHS is the result that one would obtain with the constant-price-of-risk model (18). Even if conditional betas and the factor risk premium $\lambda_t = \phi^f \omega_t^2$ are time-varying, an unconditional beta representation holds as long as the price of risk is constant. Only if the price of risk varies over time *and* is correlated with the conditional covariance of returns and the factor, the second term on the RHS introduces a wedge between the pricing implications for unconditional expected returns of the true conditional factor model and the

[7]The x_t term should not be regarded as a risk factor, though, because it is pre-determined. As a consequence, under the restrictions of the model, it plays no role in generating cross-sectional variation in expected returns.

econometrician's constant-price-of-risk model.

How big could $\mathrm{Cov}\{\mathrm{Cov}_t(r_{it+1}, f_{t+1}), \phi_t^f\}$ be for a value minus growth portfolio? Consider the case of the conditional CAPM where f is a market index log return. To maximize time-variation in the price of risk, consider the case in which all time-variation in the market risk premium is driven by time-variation in ϕ_t^f rather than by time-varying conditional factor volatility, so that $\lambda_t = \phi_t^f \omega^2$.[8] With $\beta_{it} = \mathrm{Cov}_t(r_{it+1}, f_{t+1})/\omega^2$ we obtain $\mathrm{Cov}\{\mathrm{Cov}_t(r_{it+1}, f_{t+1}), \phi_t^f\} = \mathrm{Cov}(\beta_{it}, \lambda_t)$. To get an upper bound, consider the case in which β_{it} and λ_t are perfectly correlated. As LN point out, plausible values for the standard deviation of β_{it} for a value-growth spread portfolio are around 0.30. Predictive regressions or calibrated models of time-varying risk premia like Campbell and Cochrane (1999) suggest a standard deviation of λ_t of at most 0.5% monthly. Thus, an upper bound for the wedge amounts to 0.15% monthly, or less than 2% in annualized terms—much less than the empirically observed value-growth spread of 0.50% monthly. If the correlation of β_{it} and λ_t takes a more realistic value substantially closer to zero, the gap to the observed value premium becomes even bigger. Therefore, if the constant-price-of-risk version of the CAPM of the form (18) fails to explain the value premium, it is not plausible that the conditional CAPM can explain it. A similar logic applies in the case with a non-traded factor, as in the consumption CAPM, but there it is the covariance between conditional consumption betas and the consumption risk premium that determines the pricing error.

This reasoning leads—seemingly—to a puzzle. Studies like Jagannathan and Wang (1996) and Lettau and Ludvigson (2001) use an SDF of the form (19) and find a dramatic improvement in explanatory power for the cross-section of stock returns compared with an unconditional one-factor model. The key to reconciling these findings is to note that LN's analysis takes into account the models' implications for the *time-series* dynamics of conditional moments: according to theory, for a significant wedge to exist between the unconditional

[8] Evidence from ICAPM and long-run risks models reviewed below in Section 4 suggest that conditional factor volatility drives at least part of time-variation in factor risk premia. This would imply lower time-variation in ϕ_t^f and smaller $\mathrm{Cov}\{\mathrm{Cov}_t(r_{it+1}, f_{t+1}), \phi_t^f\}$.

expected return predictions of the conditional factor model and its constant-price-of-risk version, $\mathrm{Cov}(\beta_{it}, \lambda_t)$ must be big. In contrast, if the model is fit to a cross-section of expected returns, without taking into account conditional moment restrictions, this restriction is ignored.[9] As a consequence, the cross-sectional estimates of ϕ_t^0 and ϕ_t^f can imply dynamics of conditional moments, via (13) and (14), that are completely out of line with the data. Nagel and Singleton (2011) illustrate this gap between implied and empirical dynamics of conditional moments in the case of the conditional consumption CAPM. They show that the cross-sectional estimates of ϕ_t^f can be more than an order of magnitude too volatile compared with models that feature strong, but plausible, time-variation in risk premia, like Campbell and Cochrane (1999).

3.4 Pitfalls with Low-Dimensional Factor Structures in Test Asset Returns

There is a further rather mechanical reason why fitting the unconditional versions (19) of conditional factor models—or in fact any multi-factor model—to the cross-section of expected returns yields small pricing errors for certain types of test portfolios. Most studies that test models of the form (19) use size and B/M portfolios as test assets. For these portfolios, three factors explain almost all the time-series and cross-sectional variation in returns: the market factor, SMB, and HML (Fama and French (1993)). Thus, within the log-normal framework here, we can write the log excess returns, $r_{t+1}^e \equiv r_{t+1} - r_t^f + \frac{1}{2}E[\sigma_{it}^2]$, on N test assets approximately as

$$r_t^e = Bp_t^e + e_t, \tag{22}$$

where B is a $N \times K$ factor loading matrix, p_t^e is a K-dimensional vector of log excess returns on the factors underlying the test asset returns (the Fama-French factors, with $K = 3$, in the case of the size- and B/M-sorted test portfolios), and the residual e_t is uncorrelated with p_t^e. Now consider the application of a model like (19), which we can write as $m_{t+1} = \phi^0 - \phi' \tilde{f}_{t+1}$,

[9] In a finite sample, by estimating $E[M_{t+1}R_{t+1}] = 1$ via GMM, or in a two-pass regression where, in the second stage, average returns of the test portfolios are regressed on their factor betas.

where $\tilde{f}_t$ is uncorrelated with e_t and has dimension $J = 3$. The $N \times J$ matrix of betas with respect to these proposed factors $\tilde{f}$ is

$$\beta = B\text{Cov}(p^e_{t+1}, \tilde{f}'_{t+1})\Omega^{-1} \tag{23}$$

where Ω is the covariance matrix of $\tilde{f}$. Note that the betas on the proposed factors are $K = J$ linear combinations of the factor loadings B. As a consequence, β explains the cross-section of $E[r^e_t] = BE[p^e_t]$ as well as B, namely perfectly, irrespective of the nature of the proposed factors p^e.[10] That β "explains" the cross-section of expected returns therefore has little economic content. It just re-discovers the fact that the test asset returns have a low-dimensional factor structure and that the proposed factors f are uncorrelated with residuals e. This is the point made in Lewellen, Nagel, and Shanken (2010) (LNS) (see also Daniel and Titman (2012)). In many existing tests of models of the type (19), the number of proposed factors is $J = 3$ and so the small pricing errors and high cross-sectional R^2 emphasized in these papers actually carry little economic meaning.

Unlike the cross-sectional R^2, goodness-of-fit specification tests, as in Hansen (1982), Shanken (1985), and Hansen and Jagannathan (1997), adjust for the low-dimensional factor structure in test asset returns. These tests are correctly sized under the null hypothesis that the model is correctly specified. However, a low-dimensional factor structure can create a power problem. When the magnitude of the pricing errors is small due to the low-dimensionality problem, these tests have little power to reject a false model. Moreover, researchers often intentionally focus on R^2-type measures of model performance rather than these specification tests to highlight the magnitude of pricing errors on "economically interesting" portfolios, rather than linear combinations of portfolios formed based on statistical considerations (Cochrane (1996)). Alas, if the test asset portfolios have a low-dimensional factor structure, R^2-type measures are not economically informative. To make the tests more

[10] As long as the correlation between the elements of p^e_t and f_t is not exactly zero. Lewellen, Nagel, and Shanken (2010) also analyze more general settings in which $J < K$ and in which f_t can be correlated with e_t.

informative, additional information needs to be brought in through judicious choice of the moment conditions that are used to estimate and evaluate the model.

3.5 More Informative Moment Conditions

LNS discuss several avenues to enhance the informativeness of the tests. For example, one can expand the set of test portfolios to include assets that do not share the low-dimensional factor structure of size and B/M portfolios. LNS add industry portfolios and find a dramatic deterioration in the cross-sectional fit of several previously proposed models of the type (19) (see, also, Daniel and Titman (2012)).

If factors of a proposed model are traded, a second possibility is to exploit the fact that theory links the prices of risk, via (13) and (14), to expected excess returns of the factors. This implies additional moment conditions for the factors themselves. Within our log-normal framework, the unconditional moment restrictions become

$$E[r^e_{t+1}] = \beta\lambda \tag{24}$$

$$E[f^e_{t+1}] = \lambda. \tag{25}$$

To what extent estimation pays attention to the additional restriction (25) depends on the estimator that is employed. A two-pass regression estimator where the second stage regresses the sample average of r^e_{t+1} and f^e_{t+1} on first-stage estimates of betas with OLS may pay little attention to (25). In contrast, maximum likelihood (ML), as in Shanken (1992), enforces (25) in terms of its sample analog by estimating $\hat{\lambda} = (1/T)\sum_{t=1}^{T} f^e_{t+1}$.[11] A similar logic applies with a GMM estimator of a linearized SDF as in (16) and moment conditions for excess returns $E[M_{t+1}(R_{t+1} - R^f_t)] = 0$. A GMM estimator with an identity weighting matrix may put very little weight on the moment condition for the factor excess return, while an efficient estimator with the optimal weighting matrix of Hansen (1982) enforces perfect pricing of the

[11] ML is asymptotically equivalent to a time-series regression of r^e_t on f^e_t, which automatically imposes $\hat{\lambda} = (1/T)\sum_{t=1}^{T} f^e_{t+1}$ (Gibbons, Ross, and Shanken (1989)), or a GLS cross-sectional regression which does the same (Shanken (1992)).

factor asymptotically.[12] Thus, inclusion of (25) among the moment conditions combined with the use of efficient estimation methods is likely to enhance the informativeness of the tests.[13]

Empirical work in cross-sectional asset pricing typically works with unconditional moment restrictions for estimation and evaluation, but *conditional* moment restrictions can enhance the informativeness of tests. This enhancement is especially relevant for conditional factor models. The time-varying prices of risk in these models imply strong predictions about the time-variation of conditional risk premia. It seems useful to ask whether these time-series predictions are in line with the data. With traded factors, Lewellen and Nagel (2006) show that this can be done by utilizing higher frequency data (e.g., daily returns), and estimating a factor model as a time-series regression within short time windows (e.g., quarters). The idea is that conditional betas and the conditional market risk premium (co-)vary very little within a short time window, so a time-series regression within a short window produces approximately unbiased estimates of conditional betas and conditional alphas. The conditional factor model can then be tested by examining whether the average of the (noisy) short-window alpha estimates is significantly different from zero. Based on this method, LN find that the conditional CAPM fails to explain the size and value premia. Ang and Kristensen (2012) refine this approach by using a more general non-parametric estimation methodology (LN's estimator is a special case of theirs) and they find similar results.

With non-traded factors, we can exploit conditional moment restrictions by specifying instruments and working with the moment conditions (2). But which instruments should we use? Even if we fix the set of random variables that we want to condition on, there is still an infinite number of ways to construct the instruments and to choose the number of moment conditions, as any function of $\mathcal{A}_t$-measurable random variables is a valid and feasible

[12] This can be seen by noting that the asymptotic weighting matrix S^{-1} is proportional to the inverse of the covariance matrix Ω of the residuals in a regression of returns on factors (see, e.g., the appendix of Jagannathan and Wang (2002)). If the factor is included among the test asset returns, the corresponding element of the residual vector has zero variance, and hence $S^{-1} \propto \Omega^{-1}$ puts infinite weight on pricing the factor perfectly.

[13] For non-traded factors, this approach is not available. One might think that replacing non-traded factors with mimicking portfolios might help, but this does not address the problem. One can show that if test asset returns follow (22), $K = J$, and f_t is uncorrelated with e_t, the mimicking portfolio always price the test assets perfectly even if the mimicking factor risk premia are restricted to be equal to their expected excess returns.

instrument. Some instruments may enhance the informativeness of the moment conditions, while others may add little information over and above the unconditional moment restrictions. Nagel and Singleton (2011) build on Hansen (1985) and Chamberlain (1987) to implement an optimal-instruments estimator that uses the statistically most informative instruments. Nagel and Singleton examine several scaled consumption CAPMs of the form (19) and find that these models are unable to fit the cross-sectional and time-series dimension of the data simultaneously (see, also, Roussanov (2011) for related work).

In summary, the conditional factor models that seem to do so well in explaining the cross-section of unconditional expected returns on size and B/M portfolios struggle to explain the data when additional moment conditions are brought in—as one would expect, given the discussion in Section 3.3. More generally, these results highlight that researchers must be careful to choose moment conditions that can be informative about the economic explanatory power of a model.

4 Recent Advances in Modeling the SDF

As the conditional factor model route proved to be less fruitful than it initially seemed, other approaches to enhance standard specifications of the SDF have recently received more attention. A number of recent papers work with the general version of the Epstein-Zin-Weil SDF (7). They obtain empirically implementable models by making additional assumptions about the observability of wealth portfolio returns (ICAPM) or about the dynamics of consumption (long-run risks). I also review work that links asset prices to the first-order conditions of producers, of financial market participants, and of financial intermediaries.

4.1 Substituting Out Consumption Using the Intertemporal Budget Constraint

The Epstein-Zin-Weil SDF (7) features a consumption growth factor. As consumption is notoriously difficult to measure, replacing consumption growth with easier-to-observe variables could perhaps improve the empirical performance of the model. Campbell (1993) shows that the intertemporal budget constraint can be used to substitute out consumption. Recent work by Campbell, Giglio, Polk, and Turley (2012) extends this approach to a setting with conditional heteroskedasticity. The intertemporal budget constraint implies that the log return on wealth follows

$$r^w_{t+1} = -a_t + \Delta c_{t+1} + v_{t+1}, \tag{26}$$

where $a_t \equiv \log((W_t - C_t)/C_t)$ and $v_{t+1} \equiv \log(W_{t+1}/C_{t+1})$. Campbell et al. further approximate $v_{t+1} \approx \kappa + \rho a_{t+1}$, where ρ takes a value close to one. Using this approximation in (26) and substituting into the log Epstein-Zin-Weil SDF (7) yields

$$m_{t+1} = \theta \log \delta + \frac{\theta}{\psi}\kappa - \frac{\theta}{\psi}a_t + \frac{\theta}{\psi}\rho a_{t+1} - \gamma r^w_{t+1}. \tag{27}$$

Assuming that asset returns and consumption growth are conditionally jointly log-normal and applying the conditional pricing restriction $E_t[M_{t+1}R_{t+1}] = 1$ with the SDF (27) to the return on wealth allows solving forward for a_t. Substituting this result for a_{t+1} in (27) and absorbing the conditional mean of the SDF in g_t, we obtain

$$m_{t+1} = g_t - \gamma N^{cf}_{t+1} + N^{dr}_{t+1} + \frac{1}{2}N^{risk}_{t+1} \tag{28}$$

where

$$N^{dr}_{t+1} \equiv (E_{t+1} - E_t)\sum_{j=1}^{\infty} \rho^j r^w_{t+1+j}, \tag{29}$$

$$N^{cf}_{t+1} \equiv r^w_{t+1} - E_t[r^w_{t+1}] + N^{dr}_{t+1}, \tag{30}$$

$$N^{risk}_{t+1} \equiv (E_{t+1} - E_t)\sum_{j=1}^{\infty} \rho^j \mathrm{Var}_{t+j}(m_{t+1+j} + r^w_{t+1+j}). \tag{31}$$

Thus, one obtains a conditional three-factor model that represents a discrete-time version of the intertemporal CAPM (ICAPM) of Merton (1973). The N^{dr} and N^{risk} factors reflect the exposure of the SDF to changes in investment opportunities in the form of news about expected returns and news about risk, respectively. N^{cf} represents news about expected cash flows on the wealth portfolio.

Campbell and Vuolteenaho (2004) (CV) investigate the homoskedastic version of (28) where N_{t+1}^{risk} is zero, as in Campbell (1993). Empirical implementation of the model requires forecasts of returns to compute N^{dr}. CV use a vector-autoregression (VAR) to produce these forecasts. They find that the high market beta of growth stocks in the post-1963 period arises largely from a strong negative loading on discount-rate shocks, while value stocks have higher cash-flow betas (see, also, Bansal, Dittmar, and Lundblad (2005)). With $\gamma > 1$, cash-flow risks carry a higher risk premium than discount-rate risks, resulting in a value premium. When evaluating the fit of their model, CV take into account the fact that the ICAPM restricts the risk prices in (28) to be a function of only one parameter (γ). Incorporating this restriction in the empirical estimation helps to avoid the problems associated with a low-dimensional factor structure in test asset returns that I reviewed in Section 3.4.

Campbell, Giglio, Polk, and Turley (2012) investigate the full heteroskedastic version of (28) (Brennan, Wang, and Xia (2004) empirically analyze a related model). Campbell et al. assume a one-factor structure of volatility shocks. In this case the conditional variances that enter N^{risk} are proportional to the common volatility factor. In the post-1963 sample, they find that value stocks have small positive exposure to these volatility shocks, but the exposure of growth stocks is much bigger. The positive correlation of equity returns with volatility shocks is counter-intuitive, but since volatility shocks have a negative price of risk, this further helps to explain the value premium.

Overall, the ICAPM shows some promise for explaining the cross-section of expected returns, but some challenges remain. VAR dynamics are subject to substantial estimation uncertainty, and empirical results differ to some extent across plausible specifications of investment opportunity dynamics (see Chen and Zhao (2009) and Bansal, Kiku, Shaliastovich, and Yaron (2012)). A second challenge is that the dynamics of investment opportunities estimated by the VAR may not be internally consistent with the model. The Epstein-Zin-Weil SDF implies equilibrium restrictions on the joint time-series dynamics of risk premia and volatility that are ignored in an unrestricted VAR. Taking these unrestricted VAR dynamics as exogenously given, an Epstein-Zin-Weil investor would want to time the market, while the ICAPM implementation of Campbell et al. assumes that the investor holds 100% of wealth in the market portfolio of risky assets. Holding a market-timing portfolio would in turn change the investor's view about the riskiness of value and growth stocks. Bansal, Kiku, Shaliastovich, and Yaron (2012) explore a specification in which the market risk premium is linear in the conditional variance of the market return, which is a step (albeit not a complete one) towards reconciling the dynamics of investment opportunities with the model.

4.2 Substituting Out the Wealth Return by Modeling Long-Run Consumption Dynamics

While the ICAPM approach is motivated by the notion that returns on wealth may be easier to observe than consumption, one can also make the argument that the wealth portfolio return may be difficult to observe because many important components of wealth are traded in markets subject to severe frictions (e.g., housing), or not traded at all (e.g., human capital). If one follows this argument, then substituting out the wealth return rather than consumption growth could be preferable.

Using (26) to substitute out the wealth return from the log Epstein-Zin-Weil SDF (7) yields

$$m_{t+1} = \theta \log \delta + (\theta - 1)\kappa - \gamma \Delta c_{t+1} - (\theta - 1)a_t + (\theta - 1)\rho a_{t+1}. \tag{32}$$

Applying this SDF to the conditional pricing restriction $E_t[M_{t+1}R_{t+1}] = 1$ for the wealth return and substituting for the wealth return in the conditional expectation (but not the variance) using (26), one can solve forward for a_t. Substituting the result into (32) for a_{t+1}

and absorbing the conditional mean of the SDF in g_t, we obtain

$$m_{t+1} = g_t - \gamma N_{t+1}^{sc} - (\gamma - \frac{1}{\psi}) N_{t+1}^{lc} + \frac{1}{2} \frac{\gamma - \frac{1}{\psi}}{\gamma - 1} N_{t+1}^{risk}, \tag{33}$$

where

$$N_{t+1}^{sc} \equiv \Delta c_{t+1} - E_t[\Delta c_{t+1}], \tag{34}$$

$$N_{t+1}^{lc} \equiv (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j \Delta c_{t+1+j}, \tag{35}$$

and N^{risk} is defined as before in (31). Restoy and Weil (2011) derive the homoskedastic version of (33) without N^{risk}. If $\gamma > 1/\psi$, assets earn a positive risk premium for positive correlation with shocks to long-run consumption growth, N_{t+1}^{lc}. Bansal and Yaron (2004) label this risk as *long-run* risk and they propose that consumption growth contains a small highly persistent component that makes this risk relevant for asset pricing.

Consistent with a premium for long-run risk, Parker and Julliard (2005) find that the covariance of stock returns with future consumption growth several years ahead is positively related to B/M and negatively related to size, which helps explain the cross-section of average returns on size and B/M portfolios. Liew and Vassalou (2000) find that the Fama-French size and B/M factors help predict GDP growth at an annual horizon, which also points in this direction. Hansen, Heaton, and Li (2008) find similar results with a different approach to measuring long-run risk. Bansal, Kiku, and Yaron (2007) assume a parametric model for consumption growth with a small highly persistent component and stochastic volatility. They show that N^{lc} and N^{risk} can be constructed empirically by projecting consumption growth and squared consumption innovations on the market's price-dividend ratio and the risk-free rate. Their model has a somewhat harder task than the ICAPM, because it links consumption dynamics and asset returns in an internally consistent fashion. The model explains part of the value and size premia during the 1930 to 2002 period, but it leaves a substantial portion unexplained. Standard errors are big because the long-run consumption dynamics are estimated with low precision.

These problems in estimating long-run consumption dynamics raise some doubts about the merit of the rational expectations assumption in these models. Similar concerns arise in the ICAPM, where the long-run dynamics of discount rates are difficult to pin down. As Hansen, Heaton, and Li (2008) emphasize, if econometricians struggle to pin down these long-run dynamics, it is a leap of faith to endow agents with knowledge of the parameters of the data-generating process. Developing an asset-pricing framework that incorporates learning, parameter uncertainty, and model uncertainty seems particularly relevant for models whose asset-pricing implications are driven by long-run dynamics. In fact, Collin-Dufresne, Johannes, and Lochstoer (2012) show that learning could be the source of long-run risk: a Bayesian agent learning about a constant consumption growth rate is aware that her subjective beliefs about the mean growth rate evolve as a martingale, and thus perceives long-run risk.

4.3 Production-Based Approaches

While the consumption-based asset-pricing model links the SDF to consumers' marginal rate of substitution state by state, production-based models link the SDF to producers' marginal rate of transformation (Cochrane (1991)). This raises the prospect that one could recover the SDF from (possibly better measured) macroeconomic data on the production side without relying on (hard-to-measure) consumption data. Typical production functions, however, do not allow the marginal rate of transformation to vary across states, which makes recovery of the SDF impossible without relying on investors' first-order condition, too.

The model of Belo (2010) is an exception. Belo specifies a production function that allows producers to move productivity across states, and he works out how to empirically recover the SDF as a function of industry-level production variables from the producer first-order conditions, without any assumptions about consumer preferences. He finds that this SDF does quite well in pricing size, B/M, and industry portfolios.

Cochrane (1996) uses an alternative approach. If firms have a constant returns-to-scale technology, access to financial markets, and they trade to remove arbitrage opportunities between physical investment and financial markets, then the pricing equation $E_t[M_{t+1}R^I_{t+1}] = 1$ holds for physical investment returns R^I in the same way it holds for stock returns. Just as one can construct a pricing kernel by projecting M on securities returns (Hansen and Jagannathan (1991)), one can construct one by projecting M on investment returns. Restricting the number of physical investment returns to a few common factors, Cochrane arrives at an ad-hoc factor model where the SDF is linear in two investment returns, approximated by investment growth rates (private residential and non-residential investment). Li, Vassalou, and Xing (2006) use a similar approach with a more disaggregated set of investment growth rates. These ad-hoc factor models have some success in fitting the cross-section of stock returns, but their economic content is limited. A test of these models really just tests the law of one price under the inclusion of physical investment returns, combined with the ad-hoc restriction to a small number of factors.

A third approach, pioneered by Berk, Green, and Naik (1999), is to exogenously fix an SDF as a function of some aggregate shocks, and proceed to model the evolution of firms' asset composition and systematic risk. Berk et al. use a dynamic real options model in which value firms consist of assets in place and growth options, while growth stocks are pure growth options. They exogenously specify a log SDF of the form

$$m_{t+1} = -r^f_t - \frac{1}{2}\sigma^2_z - \sigma_z \nu_{t+1}, \tag{36}$$

where ν_{t+1} is also the only aggregate shock that affects firms' assets-in-place cash flows, while the value of growth options is additionally exposed to interest-rate shocks that have a component orthogonal to ν. As a consequence, value stocks have a greater negative covariance with the SDF, i.e., a higher systematic risk and expected return, than growth stocks. Carlson, Fisher, and Giammarino (2004) work with a related model with a single source of aggregate uncertainty. Zhang (2005) endogenizes firms' cash flows and systematic risk in a competitive industry equilibrium. Zhang also uses an SDF like (36), but with a time-varying loading on ν. In his model, ν is the only source of aggregate uncertainty. Novy-Marx (2011) finds support for the prediction implied by these models that operating leverage should be positively related to expected returns.

Berk et al., Carlson et al., and Zhang argue that their models provide a "rational explanation" for the value premium. This conclusion is not warranted. These models surely produce interesting insights into the origins of firms' systematic risk exposures, but these models cannot shed light on the question whether the empirically observed cross-sectional variation in expected returns reflects priced risk under rational expectations, the effects of learning, or the effects of sentiment. The exogenous specification of an SDF like (36) does not embody any economic assumptions other than no arbitrage combined with ad-hoc restrictions on the types of shocks that do and do not enter the SDF. That the no-arbitrage restriction holds does not imply that the SDF reflects "rational" pricing. For example, in an economy in which there is a single source of aggregate uncertainty, as in Zhang (2005), any systematic effects of investor sentiment would also be driven by this single aggregate shock. Consequently, under the econometrician's objective probability measure, the SDF would also take the form (36). Berk, Green, and Naik (1999) impose the ad-hoc restriction that the part of the interest-rate shock uncorrelated with ν does not show up as a shock in the SDF. But why don't investors care about that part of the shock? Are there rational reasons for them not to care about it, or are they imperfectly rational in attaching a zero price of risk to this shock?[14] Without a structural model of consumer preferences and beliefs, the model cannot answer these questions.

A natural fourth approach is therefore to jointly model consumer preferences and the production side in general equilibrium. In many cases, though, models that take this ap-

[14] A similar assumption is made in Lettau and Wachter (2007). As Lettau and Wachter and Santos and Veronesi (2010) point out, with external habit preferences like in Campbell and Cochrane (1999), the interest-rate shock would appear in the SDF and lead to a growth premium rather than a value premium.

proach also imply that a conditional consumption CAPM or even a conditional CAPM holds, which brings the models into conflict with the empirical evidence in section 3.5. For example, Gomes, Kogan, and Zhang (2003) have a single aggregate shock in their model, and the market portfolio is conditionally a perfectly correlated mimicking portfolio for consumption, and hence the conditional CAPM holds (the same happens in Zhang (2005) and Carlson, Fisher, and Giammarino (2004)). To get around these undesirable predictions, some researchers have added additional sources of uncertainty to the model. For example, Papanikolaou (2011) works with Epstein-Zin-Weil preferences and a production technology that is subject to productivity shocks and technological innovation shocks. He constructs a long-short factor portfolio that mimicks the technological innovation shocks and which helps explain the cross-section of average returns in the data.

4.4 Rare Disasters and Higher Moments

The vast majority of empirical studies in cross-sectional asset pricing work within a log-normal framework, either explicitly, by combining a log-linear SDF with the assumption of joint log-normality of returns and the SDF, or implicitly, by applying a first-order Taylor approximation to a non-linear SDF. This approach may fail to account for higher moments and tail risks that investors care about. Recent work by Barro (2006) revived interest in the pricing implications of disaster risks.

Empirical work on the cross-sectional pricing implications of disaster risks is still in its infancy. In principle, one can account for these in empirical work by applying non-linear estimation methods (e.g., GMM) to $E_t[M_{t+1}R_{t+1}] = 1$ without linearizing the SDF or assuming log-normality. Alternatively, one can include higher-order polynomial terms in the approximation. This is the approach taken by Harvey and Siddique (2000) and Dittmar (2002). To the extent that tail risks are important, sample selection and the finite-sample distribution of estimators are of particular relevance.

4.5 Participation Constraints and Liquidity

In typical empirical applications of representative-agent models, consumption is measured as the (per-capita) aggregate consumption of all consumers in the economy. In reality, not all consumers participate in financial markets. Asset prices might reflect the consumption risks of financial market participants, but not necessarily the consumption risk of non-participants. To explore this idea, Mankiw and Zeldes (1991) construct a series of per-capita *stockholder* consumption, and they show that with this series it is easier to explain the equity premium in a consumption-based asset pricing model than with overall per-capita consumption because stockholder consumption has a higher covariance with stock returns. Malloy, Moskowitz, and Vissing-Jørgensen (2009) apply this idea to the cross-section of stock returns within a long-run risks model in which the log SDF follows a homoskedastic version of (28). They construct a stockholder consumption growth series from Consumer Expenditure Survey data from 1982 to 2004. They are able to explain cross-sectional variation across size and B/M portfolios with a much lower coefficient of relative risk aversion than with standard measures of per-capita consumption growth.

Much of the trading activity in financial markets is undertaken by financial institutions, not households. The idealized notion that agents optimally, frequently, and timely make decisions according to their first-order conditions seems more likely to hold for financial institutions than for households. Transaction costs, information costs, limited attention, and behavioral biases on the part of households may reduce the informativeness of aggregate consumption data about priced risks in financial markets. It therefore seems worthwhile to explore the first-order conditions of financial institutions.

From a financial institution's viewpoint, the definition of good and bad states of the world likely depends on factors that are not well captured by aggregate consumption. One important concern for financial institutions seems to be liquidity. This reasoning motivated Holmström and Tirole (2001), Acharya and Pedersen (2005), and Brunnermeier and Pedersen (2009) to propose models in which the SDF is liquidity-state-dependent. The notions of liquidity are different in each of these models (the marginal value of liquidity services of tradeable assets in Holmstrom-Tirole; transaction costs in Acharya-Pedersen; funding liquidity in Brunnermeier-Pedersen), but a general feature is that the consumption-based SDF M^c_{t+1} is augmented with a factor Λ_t that captures the state of aggregate liquidity

$$M_{t+1} = M^c_{t+1} \frac{\Lambda_t}{\Lambda_{t+1}}. \tag{37}$$

Augmenting the standard Epstein-Zin-Weil SDF (7) in this way yields, in logs,

$$m_{t+1} = \theta \log \delta - \frac{\theta}{\psi} \Delta c_{t+1} + (\theta - 1) r_{w,t+1} - \Delta \lambda_{t+1}, \tag{38}$$

where $\Delta\lambda_{t+1} \equiv \log(\Lambda_{t+1}/\Lambda_t)$. In this model, assets that pay off well in states in which liquidity is scarce (low Λ_{t+1}) trade at a price premium and have a low expected return.

Pástor and Stambaugh (2003) construct an empirical liquidity factor by measuring the (illiquidity-driven) tendency of individual stock returns to revert at daily frequency. They find a substantial spread in average returns between stocks that load differently on their liquidity factor and a liquidity-factor-mimicking portfolio helps explain the momentum effect (see, also, Sadka (2006)). Acharya and Pedersen (2005) construct a liquidity factor by aggregating individual stocks' Amihud (2002) illiquidity ratios. **?**) propose broker-dealer leverage as a proxy for $\Delta\lambda_{t+1}$. Their rationale is that de-leveraging indicates deteriorating funding conditions and high marginal "utility" for financial intermediaries. A single-factor SDF with this variable as the only factor does remarkably well in explaining the joint cross-section of average returns on size, B/M, momentum, and bond portfolios.

Frazzini and Pedersen (2011) consider an economy in which some investors are leverage-constrained (see, also, Baker, Bradley, and Wurgler (2011)). In this model, $-\Lambda_t$ reflects the weighted-average Lagrange multiplier on the leverage constraints. In their model, all securities have the same covariance with $\Delta\lambda_{t+1}$. Applied to the CAPM special case of (38) in our log-normal framework, this yields

$$E_t[r_{it+1}] + \frac{1}{2}\sigma_{it}^2 = r_t^f + \psi_t + \beta_{it}(E_t[r_{t+1}^w] - r_t^f - \psi_t + \frac{1}{2}\omega_{it}^2) \tag{39}$$

Thus, the *zero-beta* rate $r_t^f + \psi_t$ is higher than r_t^f and the slope $E_t[r_{t+1}^w] - r_t^f - \psi_t + \frac{1}{2}\omega_{it}^2$ of the risk-return tradeoff is flattened by ψ_t, as in Black (1972).[15] Frazzini and Pedersen find support for this prediction in a variety of international markets and asset classes.

With regards to leverage constraints, it is also worthwhile recalling the discussion of the ICAPM in section 4.1. One unresolved problem in this literature is the assumption that the investor whose Euler equation is used to price assets is assumed to be always fully invested in the market portfolio, without timing the market, even though (exogenously specified) investment opportunities are time-varying. Campbell, Giglio, Polk, and Turley (2012) point out that one potential resolution could be to interpret the Euler equation as the first-order condition of a leverage-constrained investor. The discussion here shows that this would likely entail an additional factor in the SDF driven by the Lagrange multiplier on the constraint.

5 Sentiment

I now turn to empirical work on cross-sectional asset pricing that departs from the rational expectations and Bayesian learning paradigms. Differences in opinion,[16] excessive extrapolation of growth, underreaction, limited attention are examples of such departures. Common to these approaches is that the SDF M_{t+1} that reflects sentiment-driven investors' marginal utilities prices assets under *subjective* beliefs, i.e., $E_t^S[M_{t+1}R_{t+1}] = 1$. Under *objective* beliefs, the pricing restriction $E_t[\tilde{M}_{t+1}R_{t+1}] = 1$ holds with

$$\tilde{M}_{t+1} = M_{t+1}\frac{\xi_{t+1}}{\xi_t}, \tag{40}$$

[15]There is a wedge between the zero-beta rate and the risk-free rate in this model, because leverage-constrained investors refrain from holding the risk-free asset and it is priced by unconstrained investors.

[16]Differences in opinion could also arise under Bayesian learning with non-common priors

where ξ_t is the Radon-Nikodym derivative of the subjective probability measure with respect to the objective one. Thus, for the econometrician—who works with data generated under the objective measure—the challenge is still the same as in rational expectations theories: specify and test specifications of the SDF. The only difference is that with sentiment effects present, the econometrician has to entertain the possibility that subjective beliefs enter $\tilde{M}_{t+1}$ through ξ_{t+1}/ξ_t. These relationships between pricing under the subjective and objective measure have important implications for the interpretation of empirical tests.

Part of the empirical literature has focused on ad-hoc factor models (as in the debate between Daniel and Titman (1997) and Davis, Fama, and French (2000) about "characteristics" versus "covariances") or on theoretical models with reduced-form SDFs (see the discussion in Section 4.3) to discriminate between "risk-based" and sentiment explanations of cross-sectional return predictability. Yet, without taking a stand on some theoretical restrictions on M, these approaches cannot shed light on this question. For example, suppose that $m_{t+1} \equiv \log(M_{t+1})$ follows a linear one-factor model, and $\log(\xi_{t+1}/\xi_t)$ has one-factor representation where the single factor reflects aggregate sentiment fluctuations. In this case $\tilde{m}_{t+1} \equiv \log(\tilde{M}_{t+1})$ follows a linear two-factor model. As a consequence, it is possible that an ad-hoc factor model perfectly explains expected returns, even though pricing is influenced by sentiment.

Furthermore, risk- and sentiment-based explanations could be true simultaneously. Consider an economy with two investors. Investor A holds subjective beliefs. A's first-order condition implies that $E_t^S[M_{t+1}^A R_{t+1}] = 1$ holds under A's subjective beliefs where M_{t+1}^A reflects the risk exposures taken on by A (e.g., high exposure to assets that A is optimistic about). Investor B is a sophisticated "arbitrageur" with rational expectations, and $E_t[M_{t+1}^B R_{t+1}] = 1$ holds under B's objective beliefs where M_{t+1}^B reflects the risk exposures taken on by B (e.g., low exposure to assets that A is optimistic about). An econometrician's examination of the explanatory power of M_{t+1}^A would reveal that it misprices assets under the objective measure, suggestive of a sentiment explanation. An examination of the explanatory power of M_{t+1}^B, in contrast, would reveal that it prices assets correctly under the objective measure, consistent with a risk-based explanation. Of course, both the sentiment and risk-based explanations are true simultaneously in this example as both investors' first-order conditions hold.

5.1 Limits to Arbitrage

For sentiment to have price impact, some "limits to arbitrage" must exist that prevent sophisticated investors from completely neutralizing the effect of sentiment-driven investors.[17]

If sentiment is correlated across investors, and if it affects groups of stocks or an entire asset class in the same direction, there is a natural explanation: trading against sentiment requires taking risky positions with undiversifiable risk. Sophisticated investors want to be compensated for taking on such risk, which limits their willingness to trade aggressively against sentiment. Many cross-sectional return predictability patterns offer fairly high Sharpe Ratios, though, so this is unlikely to be a full explanation.

Frictions could constitute an additional impediment. Pontiff (1996) shows that the market values of closed-end fund shares are more likely to deviate from their net asset value if an arbitrage strategy is subject to greater residual risk and to greater costs for initiating and holding the position. In a similar vein, a number of recent papers find that various cross-sectional return predictability patterns are stronger among stocks subject to such frictions.

Short-sale constraints are one example. They are most likely to be relevant among stocks with low institutional ownership, because institutional investors, and especially passive ones, are the most active lenders of stocks to short sellers. Based on this idea, Nagel (2005) shows that cross-sectional return predictability based on several predictors, including B/M, is stronger among stocks with low institutional ownership (controlling for size).

Small stocks and stocks without analyst coverage are typically not attractive investments for professional investors with big portfolios. Griffin and Lemmon (2002) find that the value premium is bigger among these stocks, and Hong, Lim, and Stein (2000) report stronger

[17] Gromb and Vayanos (2010) survey the theoretical literature on limits to arbitrage.

momentum effects. Campbell, Hilscher, and Szilagyi (2008) find that the distress anomaly is stronger for stocks with low analyst coverage, institutional ownership, and liquidity.

Cross-sectional predictability is typically stronger among stocks with high idiosyncratic volatility; see for example Ali, Hwang, and Trombley (2003) for B/M, Lipson, Mortal, and Schill (2011) and Lam and Wei (2011) for the asset growth anomaly, Mendenhall (2004) for the post-earnings announcement drift, and Mashruwala, Rajgopal, and Shevlin (2006) for the accrual anomaly. These results could indicate that idiosyncratic risk acts as a deterrent to more aggressive exploitation of predictability by sophisticated investors (see, e.g., Pontiff (2006)), but it is not clear whether this theory is quantitatively plausible. After all, strategies based on these predictors offer relatively high Sharpe Ratios.

5.2 Predicting Returns with Cross-Sectional Sentiment Proxies

Some researchers have developed proxy measures for investors' subjective beliefs, with the goal of using these to predict returns in the cross-section.

Diether, Malloy, and Scherbina (2002) and Chen, Hong, and Stein (2002) build on the idea in Miller (1977) that if investors with divergent beliefs face short-sale constraints, prices are set by optimists, while pessimists sit on the sidelines with zero positions. Greater divergence of beliefs then implies greater overpricing and lower future returns. Diether et al. proxy for investors' belief dispersion with analyst forecast dispersion, while Chen et al. proxy for it with the (lack of) breadth of mutual fund ownership of a stock. Both find support for the prediction that expected returns are low when dispersion is high.

La Porta (1996) uses analyst forecasts as a sentiment proxy and finds that stocks with high forecasted growth perform poorly subsequently, especially around future earnings announcements, supporting the hypothesis of Lakonishok, Shleifer, and Vishny (1994) that the value premium could arise from excessive extrapolation of firms' past performance.

Teo and Woo (2004) proxy for investor sentiment with the returns of mutual funds in small/large and growth/value styles. They find that high style returns in the past years predict low returns for stocks associated with these styles, consistent with the model of Barberis and Shleifer (2003). Frazzini and Lamont (2008) identify stocks with positive sentiment as those that are held by mutual funds that have recently received a lot of inflows. They show that these stocks perform poorly in the future. Furthermore, positive-sentiment stocks tend to be growth stocks and stocks with high equity issuance activity.

5.3 Cross-Sectional Implications of Time-Varying Aggregate Sentiment

Fluctuations in aggregate investor sentiment may have cross-sectional implications if different types of stocks have different sensitivity to variations in aggregate sentiment.

Lemmon and Portniaguina (2006) find that expected returns of small stocks and stocks with low institutional ownership are low when the University of Michigan consumer sentiment index is high. Baker and Wurgler (2006) construct an index from several aggregate sentiment proxies from the prior literature. They find that expected returns of stocks that are unattractive to arbitrageurs (young, small, unprofitable, extreme-growth, distressed, and high volatility stocks) are most sensitive to sentiment. Stambaugh, Yu, and Yuan (2012) use the Baker-Wurgler index to examine time-variation in a broad set of existing cross-sectional predictability patterns (most of those discussed in sections 2.2 and 2.3). They find that long-short strategies constructed from these predictors have low returns following high aggregate sentiment, and most of this time-variation comes from the short component of the portfolio, consistent with a short-sale constraints explanation. It is noteworthy that none of these papers finds that the value premium can be predicted with aggregate sentiment proxies. An exception is Ben-Rephael, Kandel, and Wohl (2012). They construct an aggregate sentiment index from a measure of investor reallocation between equity and bond funds and find that high sentiment predicts low returns most strongly among small stocks and growth stocks.

The above studies focus on directional measures of sentiment, but some researchers have also explored aggregate measures of belief dispersion. Cross-sectional implications arise if belief dispersion for individual stocks is highly sensitive to fluctuations in aggregate belief dispersion for some stocks, and less for others, or if some stocks are more likely to be affected by short-sale constraints than others. Yu (2011) constructs an aggregate measure of belief dispersion from individual stock analyst forecasts and finds that high dispersion forecasts a strong value premium. Hong and Sraer (2012) argue that high-beta stocks have the highest exposure to aggregate belief dispersion. As a consequence, a positive relationship between risk (beta) and return arises only in times of low belief dispersion. In times of high dispersion, this relationship follows an inverted U-shape. They find empirical support for these predictions. This theory is an alternative (based on heterogeneity in beliefs) to the theory of Frazzini and Pedersen (2011) (based on heterogeneity in risk aversion and leverage constraints) to explain the low average returns of high beta stocks discussed in section 4.5.

5.4 Limited Attention

A different line of work examines whether cross-sectional return predictability could arise from investors failure to pay timely attention to valuation-relevant pieces of information. Sims (2003) suggested limited attention as an explanation of inertia in macroeconomic variables, and Peng and Xiong (2006) and Hirshleifer, Lim, and Teoh (2011) apply these ideas in models of cross-sectional asset pricing. Sloan (1996) and Hirshleifer, Hou, Teoh, and Zhang (2004)) suggest limited attention as an explanation of the accrual anomaly: investors seem to pay insufficient attention to the heterogeneity in the time-series properties of different earnings components.

Evidence in a number of empirical studies is consistent with investors paying limited attention to relatively complex information. DellaVigna and Pollet (2007) show that demographic information can predict industry profitability far in the future, but the market seems to be slow in incorporating this information, resulting in industry-return predictability. Cohen and Frazzini (2008) examine economically linked firms and find that the market does not fully incorporate these economic links when responding to news. Following a similar logic, Menzly and Ozbas (2010) find cross-industry return predictability for economically linked industries. Cohen and Lou (2011) empirically categorize firms into complicated and easy-to-analyze ones and find that the returns of easy-to-analyze firms predict the returns of their more complicated peers. Belo, Gala, and Li (2012) document predictable variation in cash-flows and expected returns related to political cycles.

Hou, Peng, and Xiong (2009) look at aggregate proxies of investor attention. They argue that high trading volume and being in an up-market proxies for high investor attention, and they find that earnings momentum (interpreted as underreaction) weakens with investor attention while price momentum (interpreted as overreaction) strengthens.

6 Out-of-Sample Tests

The asset-pricing approaches that I have discussed so far may be too stark in their treatment of beliefs. Rational expectations theories assume that investors act in equilibrium as if they knew the specification and parameters of the equilibrium law of motion that generates the data. This may be a useful benchmark, but real-world investors have to learn the data-generating process. To the extent regimes change, parameters drift, and memory may get lost, it is not clear how fast, if at all, this eventual convergence to rational expectations happens. Sentiment theories on the other hand rest on the notion that investors have some fixed biases and attention deficits. Real-world investors may indeed have biases and attention deficits, but they are probably not be completely immune to learning either.

If investors learn and adapt, this has implications for cross-sectional return predictability (see Lo (2004) for an exposition of this view). There are still many open questions about how learning affects asset prices, and the body of empirical literature that deals with these issues is still quite small. Nevertheless, some progress has been made. Much of this work employs out-of-sample testing methods.

Out-of-sample tests can also help address data-snooping concerns. I review empirical work that speaks to the data-snooping issue after the discussion of learning.

6.1 Learning and Out-of-Sample Evaluation

Intuitively, one might think that if learning is driving return predictability, this predictability should not exist out of sample (where "out of sample" here is meant to be forward in time, not backwards or cross-country). But it is not as simple as that. As an illustration of the problem, consider the model of Lewellen and Shanken (2002) (LS). LS set up a model in which a risk-neutral Bayesian investor observes dividends drawn from an IID normal distribution with mean ρ. The investor initially starts with a diffuse prior. The gross risk-free rate is constant at R^f. After observing t observations of dividends, the investor's posterior mean is the sample mean $\bar{D}_t$. Now define the excess dollar return $R^e_{t+1} \equiv P_{t+1} + D_{t+1} - R^f P_t$. Combining this definition with the valuation relation $P_t = \bar{D}_t/(R^f - 1)$ yields

$$R^e_{t+1} = \left(1 + \frac{1}{(t+1)(R^f - 1)}\right)(D_{t+1} - \bar{D}_t). \tag{41}$$

From the investor's viewpoint, $E^B_t[D_{t+1}] = \bar{D}_t$ and so $E^B_t[R^e_{t+1}] = 0$, i.e., excess returns are completely unforecastable in real time. However, an econometrician examining a trading strategy that buys $\bar{D}_t$ units of the stock at t, financed at R^f, finds that it has an expected return of

$$E[R^e_{t+1}\bar{D}_t] = -\left(1 + \frac{1}{(t+1)(R_f - 1)}\right)\frac{\sigma^2}{t} < 0 \tag{42}$$

This trading strategy uses only information available in real time, and hence it is implementable out of sample, yet it earns non-zero expected returns. The reason is that the econometrician's evaluation of the population moment in (42) (or its large sample approximation) uses information about ρ that the investor does not have in real time (at time t the investor's "best guess" of ρ is $\bar{D}_t$). Lack of out-of-sample forecastability is therefore not generally a suitable null hypothesis for testing whether return predictability is generated by learning.

As the expression (42) shows, the magnitude of the trading strategy's expected excess return declines with the number of periods that the investor has learned the dividend process.

This time-decay could potentially provide a signature of learning that one could look for in the data. The form of the time-decay will typically depend on the specifics of the process that the investors are learning about, but the prediction that return predictability in out-of-sample tests is weaker than in original in-sample tests may be a robust one.

The learning problem faced by real-world investors may be of a somewhat different nature, though. In Bayesian learning models, the agent can costlessly pay attention to any number of variables that might be relevant (within the set of models entertained by the agent). In practice, even sophisticated investors and their computer algorithms can pay attention only selectively to a limited set of variables. Over time they may discover new ones, and they form an opinion about their relevance. Schwartzstein (2012) provides a model that formalizes this type of learning under selective attention. In his model, an agent pays attention to a predictor z only with a certain probability that depends on the agents belief that z is an important predictor. Eventually, as the agent updates the belief on the importance of z, the agent will learn to pay attention to z, but the learning process may take a long time. Applied to cross-sectional asset-pricing, learning under selective attention would mean that (abnormal) return predictability associated with a predictor z can exist as long as investors pay insufficient attention to z. Once attention picks up (for example, when an academic study publicizes the predictive power of z), investors quickly incorporate z into their return-forecasting models and the predictability associated with z disappears. Out-of-sample tests that check for post-publication predictability can shed light on this type of learning.

A number of studies have examined whether the evidence for cross-sectional return predictability is weaker out of sample, subsequent to the publication of academic studies. Dimson and Marsh (1999) find that the small-firm premium in the UK was not evident from 1989-1997 after studies publicized the UK small firm premium. Schwert (2003) finds that the size and value effects in the US disappeared during the 1994-2002 period. Green, Hand, and Soliman (2011) find that the accrual anomaly has disappeared post publication. In contrast, Jegadeesh and Titman (2001) find that momentum profits continued to exist in the years following the publication of Jegadeesh and Titman (1993).

One problem with these analyses is that the out-of-sample period is short, so it is difficult to distinguish true disappearance of predictability from disappearance by chance. One possible remedy is to consider many return predictors simultaneously. McLean and Pontiff (2012) study the post-publication return predictability associated with 82 characteristics that academic studies had identified. They estimate that the average decay in return predictability is about 35%.

The studies above are true out-of-sample studies. They use data samples that were not yet available to the researchers that published the original research. Decay in predictability can also be studied with *pseudo* out-of-sample tests in which the researcher splits a data sample into a training sample (used to estimate a predictive relationship) and a forecasting sample (used to evaluate the predictive performance of the forecasting model). Haugen and Baker (1996) use this approach and find substantial pseudo out-of-sample predictability for a variety of predictors. Lewellen (2011) constructs a summary measure for expected returns from 15 predictors. Using recursively expanding windows as training samples and a one-month ahead horizon for forecasting samples, he finds a decay in predictive ability of about 20-30% compared with in-sample estimates. Cooper, Gutierrez, and Marcum (2005) find only small out-of-sample predictability with a strategy that picks the best performing size, B/M, and momentum-sorted portfolios in 10-year training samples. Chordia, Subrahmanyam, and Qing (2011) find that momentum and turnover-related predictability is weaker among liquid stocks in the second half of the 1976-2009 period.

Overall, the evidence indicates that predictability tends to be somewhat weaker out of sample, but, for most cross-sectional predictors, the predictability is still substantial, even out of sample. To what extent the magnitude of the decay is consistent with learning is still an open question.

6.2 Data Snooping and Out-of-Sample Evaluation

Weaker predictability in true out-of-sample tests can also be an indication of data-snooping problems in the original studies. Correlations between a predictor z found in sample could reflect an in-sample correlation between z and noise rather than correlation of z with the true predictable component of returns. If econometricians undertake specification searches looking for predictors with significant predictive power, conventional procedures of statistical inference do not account for this specification search and hence overstate the statistical significance. One approach to this problem is to adjust critical values of test statistics for the specification search, as proposed in Lo and MacKinlay (1990) and White (2000), but implementing these methods may be difficult when specification searches occur in an informal manner or through sequential investigations by different researchers.

True out-of-sample tests can help address data-snooping concerns. If researchers found an in-sample relationship that happens to be just an in-sample correlation between a predictor z and noise, there is no reason for the predictive relationship to exist in a data sample that was not available to the researchers that published the initial results. In contrast, pseudo out-of-sample tests are not a solution to the data-snooping problem, because researchers can mine the data for significant pseudo out-of-sample predictability in the same way as it can be mined for in-sample predictability.

In addition to the post-publication evidence discussed above, such true out-of-sample evidence exists for a number of cross-sectional predictors. Davis, Fama, and French (2000) collect new accounting data and to construct B/M sorted portfolios prior to the 1960s and they find a robust value premium. Dimson, Nagel, and Quigley (2003) do a similar exercise to extend UK data back into the 1950s and they find a strong value premium. Other studies have examined data from non-U.S. stock markets and confirmed the existence of a value premium (Fama and French (1998)), momentum (Rouwenhorst (1998)), the idiosyncratic volatility effect (Ang, Hodrick, Xing, and Zhang (2009)), the equity issuance effect (McLean, Pontiff, and Watanabe (2009)), the asset growth anomaly (Watanabe, Yu, and Xu (2011)) and value and momentum effects in various asset classes internationally (Asness, Moskowitz, and Pedersen (2012)). In contrast, Leippold and Lohre (2012) find that the accrual anomaly is not a globally robust feature of stock returns. Thus, with few exceptions, the evidence indicates that cross-sectional return predictability uncovered in earlier studies is unlikely to be simply the result of data snooping.

7 Concluding Remarks

I conclude by highlighting a few key issues that emerged in this survey.

1. A substantial body of evidence has now moved us beyond Fama and French (1996). Empirical studies have found a number of robust cross-sectional return predictability patterns that are not subsumed by size and B/M and by the Fama-French three-factor model. For researchers using ad-hoc factor models as benchmarks in abnormal return calculations, there is now little empirical justification for relying on the Fama-French factors, rather than a factor model that incorporates these other sources of cross-sectional return predictability as well.

2. Ad-hoc factor models are useful tools for summarizing the main dimensions of cross-sectional return predictability, but they cannot be used to test risk-based rational expectations explanations of this predictability against sentiment explanations. Without restrictions on investor preferences, and hence on the risks that are relevant to investors, empirical tests cannot discriminate between these competing explanations. The same comment applies to production-based models with an exogenously specified reduced-form SDF. Recent work on the ICAPM and long-run risk models that employs such restrictions shows more promise to make progress on this question than the ad-hoc factor model literature.

3. The view of sentiment- and risk-based explanations as a dichotomy may be false to begin with. First, in a market in which sentiment-driven investors trade with rational

investors, systematic sentiment can affect prices, yet rational investors' first-order condition holds as well. Thus, a sentiment- and a risk-based explanation of cross-sectional return predictability can be true simultaneously. It is an interesting question whether recent empirical successes of the ICAPM could perhaps be interpreted in this way. Second, theories of learning and adaptation fall neither into the rational expectations nor the sentiment category. As real-world investors learn from data in similar ways as econometricians do, these theories may deserve greater attention in the study of cross-sectional asset pricing.

4. The choice of moment conditions is crucial for econometric evaluation to be informative about the performance of a model. In particular, fitting SDFs with time-varying prices of risk to the cross-section of expected returns only, without attention to conditional moment restrictions, can produce implied time-variation in conditional moments that is grossly inconsistent with the data.

References

Acharya, V., and L. H. Pedersen. 2005. Asset Pricing with Liquidity Risk. *Journal of Financial Economics* 77:374–410.

Ali, A., L. S. Hwang, and M. A. Trombley. 2003. Arbitrage Risk and the Book-to Market Anomaly. *Journal of Financial Economics* 69:355–373.

Amihud, Y. 2002. Illiquidity and Stock Returns: Cross-Section and Time-Series Effects. *Journal of Financial Markets* 5:31–56.

Anderson, C. W., and L. Garcia-Feijóo. 2006. Empirical Evidence on Capital Investment, Growth Options and Security Returns. *Journal of Finance* 61:171–194.

Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang. 2006. The Cross-Section of Volatility and Expected Returns. *Journal of Finance* 61:259–299.

Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang. 2009. High Idiosyncratic Volatility and Low Returns: International and Further U.S. Evidence*Journal of Financial Economics* 91:1–23.

Ang, A., and D. Kristensen. 2012. Testing conditional factor models*Journal of Financial Economics* 106:132 – 156.

Asness, C., T. Moskowitz, and L. H. Pedersen. 2012. Value and Momentum Everywhere. *Journal of Finance*:forthcoming.

Baker, M., B. Bradley, and J. Wurgler. 2011. Benchmarks as Limits to Arbitrage: Understanding the Low-Volatility Anomaly*Financial Analysts Journal* 67:40–54.

Baker, M., and J. Wurgler. 2006. Investor Sentiment and the Cross-Section of Stock Returns. *Journal of Finance* 61:1645–1680.

Bansal, R., R. Dittmar, and C. Lundblad. 2005. Consumption, Dividends, and the Cross Section of Equity Returns. *Journal of Finance* 60:1639–1672.

Bansal, R., D. Kiku, I. Shaliastovich, and A. Yaron. 2012. Volatility, the Macro- economy and Asset Prices. Working Paper, Duke University and University of Pennsylvania.

Bansal, R., D. Kiku, and A. Yaron. 2007. Risks for the Long Run: Estimation and Inference. Working Paper, Duke University.

Bansal, R., and A. Yaron. 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance* 59:1481–1509.

Barberis, N., and A. Shleifer. 2003. Style Investing*Journal of Financial Economics* 68:161 – 199.

Barro, R. J. 2006. Rare Disasters and Asset Markets in the Twentieth Century*Quarterly Journal of Economics* 121:823–866.

Belo, F. 2010. Production-Based Measures of Risk for Asset Pricing. *Journal of Monetary Economics* 57:146–163.

Belo, F., V. D. Gala, and J. Li. 2012. Government Spending, Political Cycles and the Cross-Section of Stock Returns. *Journal of Financial Economics*:forthcoming.

Ben-Rephael, A., S. Kandel, and A. Wohl. 2012. Measuring Investor Sentiment with Mutual Fund Flows*Journal of Financial Economics* 104:363 – 382.

Berk, J. B., R. C. Green, and V. Naik. 1999. Optimal Investment, Growth Options and Security Returns. *Journal of Finance* 54:1553–1607.

Bernard, V. L., and J. K. Thomas. 1989. Post-Earnings-Announcement Drift: Delayed Price Response or Risk Premium? *Journal of Accounting Research* 27:1–36.

Black, F. 1972. Capital Market Equilibrium with Restricted Borrowing. *Journal of Business* 45:444–455.

Bray, M. 1982. Learning, Estimation and the Stability of Rational Expectations Equilibria. *Journal of Economic Theory* 26:318–339.

Brennan, M. J., A. W. Wang, and Y. Xia. 2004. Estimation and Test of a Simple Model of Intertemporal Capital Asset Pricing. *Journal of Finance* 59:1743–1775.

Brunnermeier, M. K., and L. H. Pedersen. 2009. Market Liquidity and Funding Liquidity. *Review of Financial Studies* 22:2201–2238.

Campbell, J. Y. 1993. Intertemporal Asset Pricing Without Consumption Data. *American Economic Review* 83:487–512.

Campbell, J. Y., and J. H. Cochrane. 1999. By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior*Journal of Political Economy* 107:205–251.

Campbell, J. Y., S. Giglio, C. Polk, and R. Turley. 2012. An Intertemporal CAPM with Stochastic Volatility. Working Paper, Harvard University.

Campbell, J. Y., J. Hilscher, and J. Szilagyi. 2008. In Search of Distress Risk. *Journal of Finance* 63:2899–2939.

Campbell, J. Y., and R. J. Shiller. 1988. The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors. *Review of Financial Studies* 1:195–228.

Campbell, J. Y., and T. Vuolteenaho. 2004. Bad Beta, Good Beta. *American Economic Review* 94:1249–1275.

Carhart, M. M. 1997. On Persistence of Mutual Fund Performance. *Journal of Finance* 52:57–82.

Carlson, M., A. Fisher, and R. Giammarino. 2004. Corporate Investment and Asset Price Dynamics: Implications for the Cross-section of Returns*The Journal of Finance* 59:2577–2603.

Chamberlain, G. 1987. Asymptotic Efficiency in Estimation with Conditional Moment Restrictions. *Journal of Econometrics* 34:305–344.

Chan, L. K., N. Jegadeesh, and J. Lakonishok. 1996. Momentum Strategies. *Journal of Finance* 51:1681–1713.

Chen, J., H. Hong, and J. C. Stein. 2002. Breadth of Ownership and Stock Returns. *Journal of Financial Economics* 66:171–205.

Chen, L., and X. Zhao. 2009. Return Decomposition. *Review of Financial Studies* 22:5213–5249.

Chordia, T., A. Subrahmanyam, and T. Qing. 2011. Trends in the Cross-Section of Expected Returns. Working Paper, UCLA.

Cochrane, J. H. 1991. Production-Based Asset Pricing and the Link Between Stock Returns and Economic Fluctuations. *Journal of Finance* 46:209–237.

Cochrane, J. H. 1996. A Cross-Sectional Test of an Investment-Based Asset Pricing Model. *Journal of Political Economy* 104:572–621.

Cochrane, J. H., and J. Saá-Requejo. 2000. Beyond Arbitrage: Good-Deal Asset Price Bounds in Incomplete Markets*Journal of Political Economy* 108:79–119.

Cohen, L., and A. Frazzini. 2008. Economic Links and Predictable Returns. *Journal of Finance* 63:1977–2011.

Cohen, L., and D. Lou. 2011. Complicated Firms. *Journal of Financial Economics*:forthcoming.

Cohen, R. B., C. Polk, and T. Vuolteenaho. 2003. The Value Spread. *Journal of Finance* 58:609–641.

Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer. 2012. Parameter Learning in General Equilibrium: The Asset Pricing Implications. Working Paper, Columbia Business School.

Cooper, M., H. Gulen, and M. Schill. 2008. Asset Growth and the Cross-Section of Stock Returns. *Journal of Business* 63:1609–1652.

Cooper, M., R. C. Gutierrez, and B. Marcum. 2005. On the Predictability of Stock Returns in Real Time. *Journal of Business* 78:469–500.

Daniel, K., and S. Titman. 1997. Evidence on the Characteristics of Cross Sectional Variation in Stock Returns. *Journal of Finance* 52:1–33.

Daniel, K., and S. Titman. 2006. Market Reactions to Tangible and Intangible Information. *Journal of Finance* 61:1605–1643.

Daniel, K., and S. Titman. 2012. Testing Factor-Model Explanations of Market Anomalies. *Critical Finance Review* 1:103–139.

Davis, J., E. Fama, and K. French. 2000. Characteristics, Covariances, and Average Returns: 1929 to 1997. *Journal of Finance* 55:389–406.

DeBondt, W. F., and R. Thaler. 1985. Does the Stock Market Overreact? *Journal of Finance* 40:793–805.

DeBondt, W. F., and R. Thaler. 1987. Further Evidence on Investor Overreaction and Stock Market Seasonality. *Journal of Finance* 42:557–581.

DellaVigna, S., and J. M. Pollet. 2007. Demographics and Industry Returns. *American Economic Review* 97:1667–1702.

Dichev, I. D. 1998. Is the Risk of Bankruptcy a Systematic Risk? *Journal of Finance* 53:1131–1147.

Diether, K. B., C. J. Malloy, and A. Scherbina. 2002. Differences in Opinion and the Cross Section of Stock Returns. *Journal of Finance* 57:2113–2140.

Dimson, E., and P. Marsh. 1999. Winter. Murphy's Law and Market Anomalies. *Journal of Portfolio Management* 26:53–69.

Dimson, E., S. Nagel, and G. Quigley. 2003. November/December. Capturing the Value Premium in the United Kingdom. *Financial Analysts Journal* 59:35–45.

Dittmar, R. F. 2002. Nonlinear Pricing Kernels, Kurtosis Preference, and Evidence from the Cross Section of Equity Returns*The Journal of Finance* 57:pp. 369–403.

Dumas, B., A. Kurshev, and R. Uppal. 2009. Equilibrium Portfolio Strategies in the Presence of Sentiment Risk and Excess Volatility. *Journal of Finance* 64:579–629.

Epstein, L., and S. Zin. 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Theoretical Framework. *Econometrica* 57:937–969.

Epstein, L., and S. Zin. 1991. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Investigation. *Journal of Political Economy* 99:263–286.

Fairfield, P., J. S. Whisenant, and T. Yohn. 2003. Accrued Earnings and Growth: Implications for Future Profitability and Market Mispricing. *Accounting Review* 78:353–371.

Fama, E. 1970. Efficient Capital Markets: A Review of Theory and Empirical Work. *Journal of Finance* 25:383–417.

Fama, E. F., and K. R. French. 1992. The Cross-Section of Expected Stock Returns. *Journal* *of Finance* 47:427–465.

Fama, E. F., and K. R. French. 1993. Common Risk Factors in the Returns on Stocks and Bonds. *Journal of Financial Economics* 33:23–49.

Fama, E. F., and K. R. French. 1996. Mulitifactor Explanations of Asset Pricing Anomalies. *Journal of Finance* 51:55–87.

Fama, E. F., and K. R. French. 1998. Value versus Growth: The International Evidence. *Journal of Finance* 53:1975–1999.

Fama, E. F., and K. R. French. 2006. Profitability, Investment and Average Returns. *Journal of Financial Economics* 82:491–518.

Fama, E. F., and K. R. French. 2008. Average Returns, B/M and Share Issues. *Journal of Finance* 63:2971–2995.

Frazzini, A., and O. A. Lamont. 2008. Dumb Money: Mutual Fund Flows and the Cross-Section of Stock Returns. *Journal of Financial Economics* 88:299–322.

Frazzini, A., and L. H. Pedersen. 2011. Betting Against Beta. Working Paper, NYU.

Gibbons, M., S. Ross, and J. Shanken. 1989. A Test of the Efficiency of a Given Portfolio. *Econometrica* 57:1121–1152.

Giovannini, A., and P. Weil. 1989. Risk Aversion and Intertemporal Substitution. Working Paper, NBER.

Gomes, J., L. Kogan, and L. Zhang. 2003. Equilibrium Cross-Section of Returns. *Journal of Political Economy* 111:693–732.

Goyal, A. 2012. Empirical Cross-Sectional Asset Pricing: A Survey*Financial Markets and Portfolio Management* 26:338.

Green, J., J. R. M. Hand, and M. T. Soliman. 2011. Going, Going, Gone? The Apparent Demise of the Accruals Anomaly*Management Science* 57:797–816.

Griffin, J. M., and M. L. Lemmon. 2002. Book-to-Market Equity, Distress Risk and Stock Returns. *Journal of Finance* 57:2317–2336.

Gromb, D., and D. Vayanos. 2010. Limits to Arbitrage. *Annual Review of Financial Economics* 2:251–275.

Hansen, L. P. 1982. Large Sample Properties of Generalized Method of Moments Estimators. *Econometrica* 50:1029–1054.

Hansen, L. P. 1985. A Method for Calculating Bounds on the Asymptotic Covariance Matrices of Generalized Method of Moments Estimators. *Journal of Econometrics* 30:203–238.

Hansen, L. P., J. C. Heaton, and N. Li. 2008. Consumption Strikes Back? Measuring Long-Run Risk. *Journal of Political Economy* 116:260–302.

Hansen, L. P., and R. Jagannathan. 1991. Implications of Security Market Data for Models of Dynamic Economies. *Journal of Political Economy* 99:225–262.

Hansen, L. P., and R. Jagannathan. 1997. Assessing Specification Errors in Stochastic Discount Factor Models. *Journal of Finance* 52:557–590.

Hansen, L. P., and S. F. Richard. 1987. The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models. *Econometrica* 55:587–613.

Harvey, C. R., and A. Siddique. 2000. Conditional Skewness in Asset Pricing Tests*Journal of Finance* 60:1263–1295.

Haugen, R. A., and L. Baker, Nardin. 1996. Commonality in the Determinants of Expected Stock Returns. *Journal of Financial Economics* 41:401–439.

Hirshleifer, D., K. Hou, S. H. Teoh, and Y. Zhang. 2004. Do Investors Overvalue Firms with Bloated Balance Sheets. *Journal of Accounting and Economics* 38:297–331.

Hirshleifer, D., S. S. Lim, and S. H. Teoh. 2011. Limited Investor Attention and Stock

Market Misreactions to Accounting Information. *Review of Asset Pricing Studies* 1:35–73.

Holmström, B., and J. Tirole. 2001. LAPM: A Liquidity-Based Asset Pricing Model. *Journal of Finance* 56:1837–1867.

Hong, H., T. Lim, and J. C. Stein. 2000. Bad News Travels Slowly: Size, Analyst Coverage, and the Profitability of Momentum Strategies*Journal of Finance* 55:265–295.

Hong, H., and D. A. Sraer. 2012. Speculative Betas. Working Paper, Princeton University.

Hou, K., L. Peng, and W. Xiong. 2009. A Tale of Two Anomalies: The Implications of Investor Attention for Price and Earnings Momentum. Working Paper, Princeton University.

Hou, K., C. Xue, and L. Zhang. 2012. Digesting Anomalies. Working Paper, The Ohio State University.

Huang, W., Q. Liu, S. G. Rhee, and L. Zhang. 2010. Return Reversals, Idiosyncratic Risk and Expected Returns. *Review of Financial Studies* 23:147–168.

Jagannathan, R., G. Skoulakis, and Z. Wang. 2009. The Analysis of the Cross Section of Security Returns. In *Handbook of Financial Econometrics*, ed. Y. Aït-Sahalia and L. P. Hansen. Volume 2, pp. 73–134.

Jagannathan, R., and Z. Wang. 1996. The Conditional CAPM and the Cross-section of Expected Returns. *Journal of Finance* 51:3–54.

Jagannathan, R., and Z. Wang. 2002. Empirical Evaluation of Asset-Pricing Models: A Comparison of the SDF and Beta Methods*The Journal of Finance* 57:pp. 2337–2367.

Jegadeesh, N. 1990. Evidence of Predictable Behavior of Security Returns. *Journal of Finance* 45:881–898.

Jegadeesh, N., and S. Titman. 1993. Returns to Buying Winners and Selling Losers: Implications for Market Efficiency. *Journal of Finance* 48:65–91.

Jegadeesh, N., and S. Titman. 2001. Profitability of Momentum Strategies: An Evaluation of Alternative Explanations. *Journal of Finance* 56:699–720.

Jensen, M. C. 1978. Some Anomalous Evidence Regarding Market Efficiency. *Journal of Financial Economics* 6:95–101.

Kan, R., and C. Robotti. 2008. Specification Tests of Asset Pricing Models Using Excess Returnscation Tests of Asset Pricing Models Using Excess Returns. *Journal of Empirical Finance* 15:816–838.

La Porta, R. 1996. Expectations and the Cross-Section of Stock Returns*The Journal of Finance* 51:1715–1742.

Lakonishok, J., A. Shleifer, and R. W. Vishny. 1994. Contrarian Investment, Extrapolation and Risk. *Journal of Finance* 49:1541–1578.

Lam, F. E. C., and K. J. Wei. 2011. Limits-to-Arbitrage, Investment Frictions, and the Asset Growth Anomaly*Journal of Financial Economics* 102:127–149.

Lee, C. M., and B. Swaminathan. 2000. Price Momentum and Trading Volume. *Journal of Finance* 55:2017–2069.

Leippold, M., and H. Lohre. 2012. Data Snooping and the Global Accrual Anomaly. *Applied Financial Economics* 22:509–535.

Lemmon, M., and E. Portniaguina. Winter 2006. Consumer Confidence and Asset Prices: Some Empirical Evidence*Review of Financial Studies* 19:1499–1529.

Lettau, M., and S. C. Ludvigson. 2001. Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying. *Journal of Political Economy* 109:1238–1287.

Lettau, M., and J. A. Wachter. 2007. Why Is Long-Horizon Equity Less Risky? A Duration-Based Explanation of the Value Premium*Journal of Finance* 62:55–92.

Lewellen, J. 2011. The Cross Section of Expected Stock Returns. Working Paper, Dartmouth College.

Lewellen, J., and S. Nagel. 2006. The Conditional CAPM Does Not Explain Asset Pricing Anomalies. *Journal of Financial Economics* 79:289–314.

Lewellen, J., S. Nagel, and J. Shanken. 2010. A Skeptical Appraisal of Asset-Pricing Tests. *Journal of Financial Economics* 96:175–194.

Lewellen, J., and R. J. Resutek. 2012. The Predictive Power of Investments and Accruals. Working Paper, Dartmouth College.

Lewellen, J., and J. Shanken. 2002. Learning, Asset-Pricing Tests and Market Efficiency. *Journal of Finance* 57:1113–1145.

Li, E. X., D. Livdan, and L. Zhang. 2009. Anomalies. *Review of Financial Studies* 22:4301–4334.

Li, Q., M. Vassalou, and Y. Xing. 2006. Sector Investment Growth Rates and the Cross Section of Equity Returns. *Journal of Busines*:1637–1665.

Liew, J., and M. Vassalou. 2000. Can Book-to-Market, Size and Momentum be Risk Factors that Predict Economic Growth? *Journal of Financial Economics* 57:221–245.

Lipson, M. L., S. Mortal, and M. Schill. 2011. On the Scope and Drivers of the Asset Growth Effect*Journal of Financial and Quantitative Analysis* 46:1651–1682.

Lo, A. W. 2004. The Adaptive Markets Hypothesis. *Journal of Portfolio Management* 30:15–29.

Lo, A. W., and A. C. MacKinlay. 1990. Data-Snooping Biases in Tests of Financial Asset Pricing Models. *Review of Financial Studies* 3:431–467.

Lucas, R. 1978. Asset Prices in an Exchange Economy. *Econometrica* 46:1429–1445.

Malloy, C. J., T. J. Moskowitz, and A. Vissing-Jørgensen. 2009. Long-Run Stockholder Consumption Risk and Asset Returns. *Journal of Finance* 64:2427–2479.

Mankiw, N. G., and S. Zeldes. 1991. The Consumption of Stockholders and Nonstockholders*Journal of Financial Economics* 29:97–112.

Marcet, A., and J. Sargent, Thomas. 1989. Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models. *Journal of Economic Theory* 48:337–368.

Mashruwala, C., S. Rajgopal, and T. Shevlin. 2006. Why is the accrual anomaly not arbitraged away? The role of idiosyncratic risk and transaction costs. *Journal of Accounting and Economics* 42:3 – 33.

McLean, D. R., and J. Pontiff. 2012. Does Academic Research Destroy Stock Return Predictability? Working Paper, Boston College.

McLean, D. R., J. Pontiff, and A. Watanabe. 2009. Share Issuance and Cross-Sectional Returns: International Evidence. *Journal of Financial Economics* 94:1–17.

Mendenhall, R. 2004. Arbitrage risk and post-earnings announcement drift. *Journal of Business* 77:875–894.

Menzly, L., and O. Ozbas. 2010. Market Segmentation and Cross-predictability of Returns*The Journal of Finance* 65:1555–1580.

Merton, R. C. 1973. An Intertemporal Capital Asset Pricing Model*Econometrica* 41:867–887.

Merton, R. C. 1987. A Simple Model of Capital Market Equilibrium with Incomplete Information. *Journal of Finance* 42:483–510.

Miller, E. M. 1977. Risk, Uncertainty and Divergence of Opinion. *Journal of Finance* 32:1151–1168.

Muth, J. F. 1961. Rational Expectations and the Theory of Price Movements. *Econometrica* 29:1–23.

Nagel, S. 2005. Short Sales, Institutional Investors and the Cross-Section of Stock Returns. *Journal of Financial Economics* 78:277–309.

Nagel, S., and K. J. Singleton. 2011. Estimation and Evaluation of Conditional Asset

Pricing Models. *Journal of Finance* 66:873–909.

Novy-Marx, R. 2011. Operating Leverage*Review of Finance* 15:103–134.

Novy-Marx, R. 2012a. The Other Side of Value: The Gross Profitability Premium. *Journal of Financial Economics*:forthcoming.

Novy-Marx, R. 2012b. Pseudo-Predictability in Conditional Asset Pricing Tests: Explaining Anomaly Performance with Politics, the Weather, Global Warming, Sunspots, and the Stars. Working Paper, NBER.

Ohlson, J. A. 2005. On Accounting-Based Valuation Formulae. *Review of Accounting Studies* 10:323–347.

Papanikolaou, D. 2011. Investment Shocks and Asset Prices. *Journal of Political Economy* 119:639–685.

Parker, J. A., and C. Julliard. 2005. Consumption Risk and the Cross Section of Expected Returns. *Journal of Political Economy* 113:185–222.

Pástor, Ľ., and R. F. Stambaugh. 2003. Liquidity Risk and Expeced Stock Returns. *Journal of Political Economy* 111:642–685.

Peng, L., and W. Xiong. 2006. Investor Attention, Overconfidence and Category Learning*Journal of Financial Economics* 80:563 – 602.

Pontiff, J. 1996. Costly Arbitrage: Evidence from Closed-End Funds*Quarterly Journal of Economics* 111:1135–1151.

Pontiff, J. 2006. Costly Arbitrage and the Myth of Idiosyncratic Risk. *Journal of Accounting and Economics* 42:35–52.

Pontiff, J., and A. Woodgate. 2008. Share Issuance and Cross-Sectional Returns. *Journal of Finance* 63:921–945.

Restoy, F., and P. Weil. 2011. Approximate Equilibrium Asset Prices*Review of Finance* 15:1–28.

Ross, S. A. 1976. The Arbitrage Theory of Capital Asset Pricing. *Journal of Economic Theory* 13:341–60.

Roussanov, N. 2011. Composition of Wealth, Conditioning Information, and the Cross-Section of Stock Returns. Working Paper, University of Pennsylvania.

Rouwenhorst, K. G. 1998. International Momentum Strategies. *Journal of Finance* 53:267–284.

Sadka, R. 2006. Momentum and Post-Earnings-Announcement Drift Anomalies: The Role of Liquidity Risk. *Journal of Financial Economics* 80:309–349.

Santos, T., and P. Veronesi. 2010. Habit formation, the cross section of stock returns and the cash-flow risk puzzle*Journal of Financial Economics* 98:385 – 413.

Schwartzstein, J. 2012. Selective Attention and Learning. Working Paper, Dartmouth College.

Schwert, G. W. 2003. Anomalies and Market Efficiency. In *Handbook of the Economics of Finance*, ed. G. M. Constantinides, M. Harris, G. M. Stulz, Rene Constantinides, M. Harris, and R. Stulz, pp. 939–972. Amsterdam: Elsevier North Holland.

Shanken, J. 1985. Multivariate Tests of the Zero-Beta CAPM. *Journal of Financial Economics* 14:327– 348.

Shanken, J. 1992. On the Estimation of Beta-Pricing Models*Review of Financial Studies* 5:1–33.

Sims, C. A. 2003. Implications of Rational Inattention*Journal of Monetary Economics* 50:665 – 690.

Sloan, R. 1996. Do Stock Prices Fully Reflect Information in Accruals and Cash Flows About Future Earnings? *Accounting Review* 71:289–315.

Stambaugh, R. F., J. Yu, and Y. Yuan. 2012. The Short of It: Investor Sentiment and Anomalies*Journal of Financial Economics* 104:288 – 302.

Teo, M., and S.-J. Woo. 2004. Style effects in the cross-section of stock returns*Journal of Financial Economics* 74:367 – 398.

Timmermann, A., and C. W. Granger. 2004. Efficient Market Hypothesis and Forecasting. *International Journal of Forecasting* 20:15–27.

Titman, S., K. J. Wei, and F. Xie. 2004. Capital Investments and Stock Returns. *Journal of Financial and Quantitative Analysis* 39:677–700.

Vuolteenaho, T. 2000. Understanding the Aggregate Book-to-Market Ratio and Its Implications to Current Equity-Premium Expectations. Working Paper, Harvard University.

Vuolteenaho, T. 2002. What Drives Firm-Level Stock Returns? *Journal of Finance* 52:233–264.

Watanabe, A., T. Yu, and Y. Xu. 2011. The Asset Growth Effect and Market Efficiency: Insights from International Stock Markets. Working Paper, University of Alberta.

Weil, P. 1990. Nonexpected Utility in Macroeconomics. *Quarterly Journal of Economics* 105:29–42.

White, H. 2000. A Reality Check for Data Snooping*Econometrica* 68:1097–1126.

Yu, J. 2011. Disagreement and Return Predictability of Stock Pwortfolios*Journal of Financial Economics* 99:162 – 183.

Zhang, L. 2005. The Value Premium*The Journal of Finance* 60:pp. 67–103.