Accounting anomalies and fundamental analysis: An alternative view

Jonathan Lewellen*

Tuck School of Business, Dartmouth College, Hanover, NH 03755, USA

ARTICLE INFO

ABSTRACT

The literature on accounting anomalies and fundamental analysis provides important insights into the behavior of stock prices and the relation between accounting numbers and firm value. My review discusses five key topics from this literature: (1) discriminating between risk and mispricing explanations for return anomalies; (2) estimating the implied cost of capital; (3) inferring investors' perceptions of the earnings process; (4) understanding the importance of trading costs and firm size; and (5) improving the construction of characteristic-based trading strategies. My discussion highlights important challenges facing the literature and offers suggestions for improving empirical tests.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

The literature on accounting anomalies and fundamental analysis remains one of the most active areas of research in accounting and finance. At its core, the goal of the literature is to understand how accounting numbers relate to firm value and how quickly and accurately investors assess the information in financial reports. These issues have clear implications for investment management, corporate policy (including, almost certainly, real investment decisions), standard setting, and contracting.

Richardson, Tuna, and Wysocki (RTW, this issue) provide a comprehensive review of recent empirical work in this area, focusing on four main topics: (1) fundamental analysis—the relation between accounting numbers and firm value, including estimates of a firm's implied cost of capital; (2) the accrual anomaly—the negative relation between accruals and future stock returns; (3) post-earnings-announcement drift—the observation that stock prices seem to underreact to quarterly earnings news; and (4) the value anomaly—the relation between price-based valuation ratios and future stock returns. Their survey is exemplary in many ways: it provides not only a thorough review of the literature but also evidence on how academic research relates to investment practice, new results for the accrual anomaly and post-earnings-announcement drift, and thoughtful suggestions for future research.

Much of RTW's survey focuses on the use of accounting numbers to forecast a firm's earnings, cash flows, and stock returns, with the view that 'financial statements can help investors make better portfolio allocation decisions' (RTW, this issue). The notion that forecasting represents a key use of accounting numbers seems undeniable. Indeed, almost any role for accounting numbers—far beyond their use by investors—relies, implicitly or explicitly, on their association with firm performance and value. For example, using accounting numbers in compensation and debt contracts seems sensible only if those numbers are correlated, at least somewhat, with the firm's current and future cash flows.

* Corresponding author. Tel.: +1 603 646 8650.
  E-mail addresses: jonathan.lewellen@dartmouth.edu, jon.lewellen@dartmouth.edu

RTW's survey provides a good overview of many of the issues that are central to the literature, including, for example, the role of theory, the design of empirical tests, and the importance of risk and transaction costs. My goal here is to provide an alternative perspective on the literature and to comment on the empirical methods found in many studies. In particular, I discuss five main topics: (1) distinguishing between risk and mispricing explanations for return predictability; (2) estimating the implied cost of capital imbedded in stock prices; (3) inferring investors' perceptions of the earnings process from the behavior of stock returns (so-called Mishkin tests); (4) understanding the importance of trading costs and firm size for asset-pricing tests; and (5) improving the construction of trading strategies using ex ante measures of risk (an idea advocated by RTW). My discussion highlights important challenges facing the literature and offers suggestions for improving empirical tests.

2. Risk vs. mispricing

One of the literature's key results is that earnings, cash flows, accruals, valuation ratios, and asset growth predict cross-sectional variation in expected stock returns. The patterns are well-documented and seem almost certainly to be 'real,' i.e., they are highly unlikely to be the outcome of random chance or data snooping. For example, Richardson et al. (2005, Table 8) report a t-statistic of -6.38 when annual size-adjusted stock returns from 1962 to 2001 are regressed cross-sectionally on lagged accruals (as well as earnings). Viewed in isolation, this t-statistic has a p-value of 0.0000002 if the annual Fama–MacBeth slopes are normally distributed (and IID). A t-statistic of this magnitude or larger would be observed less than 1% of the time, under the null of no predictability, even if researchers ran 50,000 statistically independent predictability tests (i.e., $0.9999998^{50,000} \approx 0.99$).

The main empirical challenge, then, is to *interpret* the results, specifically, to distinguish between risk and mispricing stories for predictability. Any attempt to do so immediately confronts Fama's (1970) joint-hypothesis problem: in order to test whether returns are anomalous, we must know what expected returns should be in the absence of mispricing. This problem is especially acute for many accounting anomalies because plausible risk and mispricing explanations both exist.

Mispricing stories for accounting anomalies are only loosely grounded in theoretical models of investor behavior and rely, instead, on the generic idea that investors do not understand the properties of earnings and accruals. For post-earnings-announcement drift, the standard argument is that investors underreact to quarterly earnings announcements and, in particular, do not fully appreciate the persistence of seasonally differenced earnings (see, e.g., Bernard and Thomas, 1990). For the accrual anomaly, the simplest story is that investors fixate on bottom-line earnings and do not fully appreciate the differential reliability (or persistence) of its cash flow and accrual components (see, e.g., Sloan, 1996; Richardson et al., 2005; Dechow et al., 2008). An alternative view is that accruals correlate with the investment decisions of the firm, which, in turn, are related to valuation errors because investors do not understand how growth impacts the firm's future profitability (see, e.g., Fairfield et al., 2003; Richardson et al., 2006). Finally, for the value anomaly, the traditional argument is that investors overreact to firms' past performance, failing to fully appreciate its transitory nature (see, e.g., Lakonishok et al., 1994).

Risk-based stories for the anomalies have evolved over the last twenty years. The basic argument is simply that different types of stocks are exposed to different amounts of systematic risk and, therefore, carry different expected returns. For example, Fama and French (1992) suggest that 'distressed' value stocks might be risky because they are especially sensitive to economic conditions. Subsequent work grounds this idea more formally in the operating risks that different firms are likely to face. Berk et al. (1999), Gomes et al. (2003), Carlson et al. (2004, 2006), Zhang (2005), and Novy-Marx (2010) show that a combination of growth options, irreversible investment, and operating leverage will naturally make small, value stocks riskier than large, growth stocks. Similarly, Li et al. (2009) and Wu et al. (2010) observe that basic q-theory predicts that profitable, slowly growing firms are likely to have higher expected returns than less profitable, quickly growing firms: in equilibrium, firms with high costs of capital (or expected stock returns) tend to be more profitable to compensate for high financing costs and to grow more slowly because they have fewer positive-NPV projects.

The risk-based models above largely ignore the crucial issue of how to *measure* risk; they take the asset-pricing factor(s) as given and focus, instead, on how risk is likely to vary cross-sectionally and through time as a function of firm characteristics. A parallel stream of research takes the opposite approach: it explores new models of aggregate risk but abstracts from the inner-workings of the firm. These models come in three forms. The first group represents new empirical versions of Merton's (1973) intertemporal CAPM, in which the risk factors capture shocks to wealth and shocks to investment opportunities (Campbell and Vuolteenaho, 2004; Brennan et al., 2004; Petkova, 2006). The second group explores conditional versions of the Sharpe–Lintner or consumption CAPM using new state variables to capture how risk and expected returns change through time (Campbell and Cochrane, 1999, 2000; Lettau and Ludvigson, 2001; Duffee, 2005; Lustig and Van Nieuwerburgh, 2004; Santos and Veronesi, 2006). The third group proposes new risk factors not found in traditional models, including shocks to labor income (Jagannathan and Wang, 1996; Heaton and Lucas, 2000; Jacobs and Wang, 2004; Santos and Veronesi, 2006), growth in macroeconomic output and investment (Cochrane, 1996; Vassalou, 2003; Li et al., 2006), growth in luxury, durable, and future consumption (Ait-Sahalia et al., 2004; Bansal et al., 2005; Parker and Julliard, 2005; Yogo, 2006; Hansen et al., 2008), and shocks to aggregate liquidity (Pastor and Stambaugh, 2003; Acharya and Pedersen, 2005). None of these models have been widely adopted in the literature, but collectively they provide a good sense that asset-pricing theory continues to evolve and that we do not yet have anything close to a

complete notion of how to measure risk and expected returns. An implication is that we still do not have a perfect solution to Fama's joint-hypothesis problem.

At the same time, however, I think it is too pessimistic to conclude that we have no way of distinguishing risk from mispricing. Indeed, the literature suggests at least two possible tests that do not require a fully specified risk model:

(1) One approach is to test for predictability in short-horizon returns, for example, around future earnings announcements (see, e.g., Bernard and Thomas, 1990; Jegadeesh and Titman, 1993; Sloan, 1996; La Porta et al., 1997). The virtue of this approach is that almost any risk-based model would predict that short-horizon expected returns should be small, allowing us to largely side-step the joint-hypothesis problem. The main limitation is that announcement returns provide only a lower bound on how much of an anomaly can be explained by mispricing, since pricing errors might be corrected slowly over time rather than at easily identified announcement dates. (The flip-side is that, if announcement returns are predictable, it would be incorrect to conclude that mispricing explains the whole anomaly—unless, of course, announcement returns do represent the whole anomaly.)

(2) A second way to disentangle risk and mispricing stories is just to consider how large an anomaly is economically: is the implied Sharpe ratio plausible in an efficient market? Is the anomaly close to an arbitrage opportunity? Again, the idea is that any risk-based model, regardless of its exact nature, should limit how profitable a trading strategy can be, helping to side-step the joint-hypothesis problem. Unfortunately, the literature does not provide a precise definition of a 'reasonable' Sharpe ratio, but I would be willing to guess that some observed Sharpe ratios are likely to be viewed as anomalously high by a majority of the profession.

Richardson et al.'s (2005) results, mentioned earlier, provide a good illustration of this argument. It is well-known that the slope coefficient in a Fama–MacBeth regression equals the return on a zero-investment portfolio (Fama, 1976) and that the associated t-statistic is proportional to the portfolio's sample Sharpe ratio: $t=r/(s/T^{1/2})=T^{1/2}q$, where r is the portfolio's average excess return (i.e., the Fama–MacBeth slope), s is the portfolio's sample standard deviation, T is the length of the time series, and q is the sample Sharpe ratio ($q\equiv r/s$). The t-statistic of 6.38 (in absolute value) reported by Richardson et al. from 1962 to 2001 implies an annual Sharpe ratio of 1.01 ($6.38/40^{1/2}$). This value is over three times larger than the Sharpe ratio of the CRSP value-weighted index during the same period, 0.32, a number that is already a challenge for the asset-pricing literature to explain with plausible levels of risk aversion (see, e.g., Mehra and Prescott, 1985).

To put these ideas in more formal terms, Hansen and Jagannathan (1991) show that the maximum Sharpe ratio available in the market provides a lower bound on the volatility of the so-called stochastic discount factor (SDF).[1] In risk-based models, the SDF is proportional to the marginal utility of wealth or consumption, low in good states when investors are rich and high in bad states when investors are poor. The SDF is scaled to have a mean of $1/(1+r_f)$ and, as Hansen and Jagannathan show, a standard deviation of at least $q_{max}/(1+r_f)$, where r_f is the riskfree rate and q_{max} is the maximum-possible Sharpe ratio. Therefore, to explain the Sharpe ratio of 1.01 implied by Richardson et al.'s results, the SDF must have a volatility of at least $1.01/(1+r_f)\approx 0.95$. This means that an extra dollar of wealth must be enormously more valuable in a bad state of the world (when the SDF is high) than in a good state (when the SDF is low). For example, with a mean around 1.0 and a volatility of 0.95, the SDF might be 2.0 in a bad state and 0.5 in a good state, implying that the marginal utility of wealth is four times higher in the bad state. These facts are hard to reconcile with typical risk-based models.

RTW suggest three alternative tests to help distinguish risk from mispricing. They argue that (1) finding that returns are positive in both up and down markets can help to rule out a risk-based story since 'a risk explanation would not be consistent with a stable return stream across overall market environments' (RTW, this issue); (2) tests that explore biases in analysts' forecasts can help to confirm that errors in expectations contribute to an anomaly; and (3) finding that an anomaly diminishes over time suggests that investors have discovered and started to trade on the anomaly and, thus, that the anomaly must have reflected mispricing in the past.

My own view is that, while each of these tests is useful, none is effective in distinguishing risk from mispricing. The first test is basically another way of looking at the beta of a trading strategy and does not really confront the central problem—that we are not sure the CAPM provides a complete model of risk. The second test assumes that analysts have the same biases as other investors, which simply substitutes one joint-hypothesis problem for another. The third test assumes that risk is stable through time and that the only reason an anomaly might change during the sample is if the stock market becomes more efficient. That assumption requires a big leap of faith and, at face value, seems to be inconsistent with much evidence of time-varying risk.

3. The implied cost of capital

A second key area of research attempts to infer a firm's cost of (equity) capital from observed stock prices and forecasts of a firm's dividends and earnings. The basic idea has a long history in the literature and, in its simplest version, can be

---

[1] An SDF is a positive random variable, m, such that $E[m(1+R_i)]=1$ for all asset returns R_i. Rearranging this expression, and using the fact that $E[m]=1/(1+r_f)$ and $E[m(1+R_i)]=\text{cov}(m,R_i)+E[m]E[1+R_i]$, leads to the familiar prediction that a stock's expected excess returns are proportional to its systematic risk, $\text{cov}(m, R_i)$. Different risk models imply different SDFs. For example, the CAPM says that the SDF is a function of the market return, while the consumption CAPM predicts that it is a function of consumption.

implemented solving the Gordon growth formula, $P=D/(r-g)$, for r, yielding $r=D/P+g$, where r is the cost of capital, D is next year's expected dividend, P is the current price, and g is the perpetuity growth rate of dividends. Applied to the S&P 500, for example, the dividend–price ratio is currently around 2% (in Spring 2010) and the expected long-term growth rate might be around 6%, either extrapolating from the historical growth rate of dividends and earnings over the past 50 years or estimated as the sum of expected inflation (say 3%), expected productivity growth (2%), and expected population growth (1%). Thus, the implied cost of capital for the S&P 500 would be roughly $r=2\%+6\%=8\%$.

The Gordon growth formula can reasonably be applied to the overall market—for which the constant growth assumption may be plausible—but is less realistic for individual stocks, which may not pay dividends or might exhibit rapid short-run growth. The solution offered by the literature is to re-write the present value formula in terms of earnings, allowing for different short-run and long-run growth rates. The literature provides, in particular, two versions of earnings-based valuation models. The first is the well-known residual income model:

$$P_t = B_t + \frac{E[RI_{t+1}]}{1+r} + \frac{E[RI_{t+2}]}{(1+r)^2} + \frac{E[RI_{t+3}]}{(1+r)^3} + ...,$$ (1)

where P_t is the stock price (at date t), B_t is the book value of equity per share, r is the discount rate, and RI_t is residual income, defined as earnings per share minus rB_{t-1}. The second version focuses on the abnormal growth in earnings (AG):

$$P_t = \frac{E[EPS_{t+1}]}{r} + \frac{1}{r}\frac{E[AG_{t+2}]}{1+r} + \frac{1}{r}\frac{E[AG_{t+3}]}{(1+r)^2} + ...,$$ (2)

where EPS_t is earnings per share and AG_t is defined as $EPS_t-EPS_{t-1}-r(EPS_{t-1}-D_{t-1})$, a measure of the firm's abnormal change in earnings relative to what would be expected if the firm simply invested last year's retained earnings, $EPS_{t-1}-D_{t-1}$, at the cost of capital (see Easton, 2004; Ohlson, 2005; Ohlson and Juettner-Nauroth, 2005). The literature then infers the cost of capital, r, from observed stock prices and forecasts of future earnings and dividends, typically using analyst forecasts as a starting point to get expected earnings for the first few years.

A source of confusion in the literature stems from some ambiguity about what exactly is being measured. Specifically, the 'implied cost of capital' could be interpreted as either the discount rate investors use to value the firm—i.e., the rate of return investors require for holding the stock—or as the true expected return imbedded in stock prices (based on rational forecasts). These two concepts are the same if the market is efficient but differ if investors have biased expectations of future dividends or earnings. Thus, if we admit the possibility of stock mispricing, we need to be clear about whether we are interested in investors' expectations of cash flows and returns or in rational forecasts of cash flows and returns. Much of the literature implicitly focuses on the latter but, in many applications, investors' required returns are probably more important, as I discuss further below.

In my view, while a firm's implied cost of capital is interesting in principle, the literature faces three big hurdles that substantially limit how useful it is in practice:

(1) Estimates of the implied cost of capital for individual firms are likely to be both noisy and biased. In a sense, the central conceit of the literature is that valuation is easy: we simply have to plug analyst forecasts into one of the models, make a reasonable long-term growth assumption, and, voilà, we obtain a good estimate of a firm's fundamental value. In truth, a firm's fundamental value is hard to estimate accurately—a view supported by the vast sums spent on fundamental analysis by the investment banking and money management industries—and analyst forecasts are unlikely to be a good proxy for either the market's expectations or true expectations.

To be more concrete, consider the basic Gordon growth model: $r=D/P+g$. Even if we have a precise forecast of next year's dividend (or earnings), noise in long-term growth forecasts translates one-for-one into noise in the implied cost of capital. So, for example, if the true discount rate and error in our growth forecasts both have a cross-sectional standard deviation of 3%, then half the variation in the implied cost of capital would be real and half would be noise. And any error in the short-run cash flow forecast, or any bias in either short-run forecasts or long-term growth rates, would exacerbate the measurement problem.

Defenders of the implied cost of capital might respond in a couple of ways. One possible response would be that, even if estimates of the implied cost of capital for individual firms are noisy, the average for a diversified portfolio of stocks could still be precise. The problem, however, is that methodological choices almost certainly introduce systematic errors in the implied cost of capital that do not cancel out in portfolios. For example, the use of analyst forecasts imparts both a general upward bias in the implied cost of capital (e.g., Easton and Sommers, 2007) and, to the extent that analysts are differentially biased for different types of stocks, a correlation between the implied cost of capital and firm characteristics (see, e.g., La Porta, 1996; Dechow and Sloan, 1997; Bradshaw et al., 2001, for evidence of differential bias). Similarly, different assumptions about long-term growth rates can induce systematic or industry-specific errors in the implied cost of capital (e.g., Botosan and Plumlee, 2005).

A second possible response is that, even if the implied cost of capital is imperfect, it might still be better than alternatives such as realized returns or estimates from asset-pricing models like the CAPM or Fama–French (1993) three-factor model. Realized returns, in particular, are often dismissed by the literature as 'unreliable' (e.g., Botosan and Plumlee, 2005; Easton and Monahan, 2010). But realized returns have the crucial advantage that they provide inherently unbiased estimates of true expected returns (for the simple reason that true expected returns are *defined* as the expectation of

realized returns, conditional on information known prior to the period).[2] In many applications, that property is much more important than precision because the potential for unknown bias in the implied cost of capital means that, at a fundamental level, we can learn little from an empirical test: does a result tell us about expected returns or bias in our empirical proxy? We have no way of knowing. (As for estimates from asset-pricing models, I am not aware of any research that explores how the precision and bias in such estimates compares with the precision and bias in the implied cost of capital.)

(2) A second challenge for the literature is to validate empirically an estimate of the implied cost of capital. The most common approach in the literature is to ask whether the implied cost of capital is a good measure of a firm's true expected return. If it were, we would expect to find a slope of one when stock returns are regressed, cross-sectionally, on the implied cost of capital. In principle, we could take the idea further and ask if the implied cost of capital provides the best estimate of expected returns: does it drive out all other predictive variables in a cross-sectional regression? In practice, the literature provides little evidence that the implied cost of capital is even positively correlated with future returns (e.g., Easton and Monahan, 2005).

However, it is not at all clear that the implied cost of capital *should be* a good estimate of a stock's true expected return. One reason was given earlier: if the market is inefficient and analysts have the same biases as other investors, then the implied cost of capital would measure investors' required return for holding a stock, not the true expected return. In fact, the implied cost of capital could be an excellent estimate of the discount rate imbedded in a firm's stock (as perceived by investors) and yet show little correlation with the stock's true expected return.

A second reason the implied cost of capital could be a poor measure of a stock's true expected return is that it represents—at best, if earnings are forecasted well—the long-run internal rate of return from holding the stock, not the short-run expected return. The two are equivalent only if expected returns are the same for all horizons. Empirically, however, expected returns seem to fluctuate with a host of time-varying firm characteristics, such as a firm's past returns, book-to-market ratio, and accruals. Indeed, there are plausible scenarios in which a stock's short-run expected return moves in the opposite direction as its long-run implied cost of capital. For example, Jegadeesh and Titman (2001) find that past winner stocks have high (true) expected returns in the short run but low expected returns in the long run. Depending on where the stock is in this cycle, its short-run expected return could be high even though its long-run internal rate of return is low (e.g., the stock may be overvalued and predicted to become even more overvalued in the short run). Similarly, Cohen et al. (2009) argue that value stocks may be priced correctly—so their implied cost of capital should be close to the appropriate level—even though they have anomalously high expected returns in the short run. As with winner stocks, high short-run expected returns for value stocks are offset by low long-run expected returns.

Thus, it seems problematic to test whether an estimate of the implied cost of capital is good by correlating it with short-run expected returns. Such a test could, in principle, produce either a type 1 or type 2 error: we might find that the implied cost of capital works poorly even though it provides a good estimate of the stock's internal rate of return, or we could find that the implied cost of capital has predictive power even if it is a poor estimate of the internal rate of return. An alternative test, closer in spirit to what the implied cost of capital actually measures, would be to correlate the implied cost of capital with long-run expected returns using an approach like that of Cohen et al. (2009).

(3) A third challenge for the literature is to explain why the implied cost of capital is useful. Consider three groups who might be interested in a firm's cost of capital:

Corporate managers: It seems obvious that managers would be interested in the firm's implied cost of capital, but there are actually good reasons to question that assumption. First, we have no evidence that the implied cost of capital is estimated precisely enough to be useful to managers. Second, even if it was estimated precisely, the implied cost of capital is often appropriate for capital budgeting only if it represents the return required by investors. If stocks are mispriced and the implied cost of capital reflects true (not biased) expectations, it will typically differ from the discount rate that should be used for a project (see, e.g., Stein, 1996). In particular, the discount rate should be determined solely by the risk of the project, under some assumptions, even if the firm has access to cheap financing because the stock is overpriced. Put differently, mispricing may affect a firm's financing decisions without affecting its real investment decisions.

Investors: The implied cost of capital—i.e., the long-run internal rate of return from holding a stock—is useful for a long-term buy-and-hold investor but, for everyone else, is probably less important than the short-run expected return forecast over, say, the next month, quarter, or year. And, as I argued above, the long-run performance of a stock, as indicated by the implied cost of capital, may have little connection to a stock's short-term expected return.

Researchers: The implied cost of capital would be useful for researchers if it provided an unbiased estimate of the required return demanded by investors. If it did, we could use the implied cost of capital to study how required returns vary with a stock's risk (e.g., Botosan and Plumlee, 2005), with firm characteristics like book-to-market, accruals, or disclosure policy (e.g., Gebhardt et al., 2001), or as a function of events that are hypothesized to affect the firm's cost of capital (e.g., Hail and Leuz, 2009). Again, however, we have little evidence that the implied cost of capital provides a good measure of market

---

[2] The only way realized returns can be 'biased' is if we condition the sample on information that is not known prior to the period. For example, the market's return in years that interest rates drop provides a biased estimate of the market's expected return, since the decline in interest rates is known only after the fact. The past returns of firms announcing an acquisition today give a biased estimate of the firms' expected returns because the sample is selected based on ex post information (the announcement of an acquisition).

expectations, and any potential bias in the implied cost of capital makes it difficult to assess whether empirical results tell us about required returns or problems with the estimates. Also, to the extent that risk changes through time, the implied cost of capital should reflect the long-run riskiness of a stock, not the short-run risk level.

In sum, I believe the literature faces a variety of challenges in validating estimates of the implied cost of capital and explaining why they are useful.

4. Mishkin tests

Many papers in the literature use so-called Mishkin (1983) tests to study whether investors' perceptions of the earnings process differ from the true time-series properties. To review, the basic idea is simple. Let X_t be earnings, Z_{t-1} be any predictive variable known at the beginning of the period, and UR_t be the unexpected stock return in period t (unexpected by investors, not necessarily unexpected given the true information set). By definition, unexpected returns—if they are correctly measured—react only to new information during period t and, thus, must be uncorrelated with the portion of earnings that is anticipated by investors. This implies that, in the following regression

$$UR_t = \beta(X_t - c'Z_{t-1}) + e_t, \tag{3}$$

the slope coefficient c that maximizes the R^2 will capture investors' expectation of the earnings process; i.e., the slope coefficient that maximizes the correlation between UR_t and $UX_t \equiv X_t - c'Z_{t-1}$ must be the one that gives the best measure of unexpected earnings as perceived by investors, so $c'Z_{t-1}$ measures investors' expected earnings conditional on Z_{t-1}. The Mishkin test compares this 'perceived' slope to the slope of the actual time-series process:

$$X_t = d'Z_{t-1} + u_t. \tag{4}$$

The null of market efficiency predicts that $c=d$. (I omit intercept terms from these regressions for expositional convenience only.)

A couple of features of the Mishkin test are immediate and have been discussed elsewhere (see, e.g., Abel and Mishkin, 1983; Kraft et al., 2007). First, in the paragraph above, I did not assume that Z_{t-1} represents the full set of information available to investors; under the null of market efficiency, the slopes c and d should be the same regardless of what other (omitted) variables help to predict earnings. In that sense, the test of market efficiency remains valid even if there are correlated omitted variables. Second, the Mishkin test is conceptually identical to a standard predictability test for returns so long as β, the slope on unexpected earnings in the return regression, is not zero. Specifically, using Eq. (4), we can re-write the return regression as follows:

$$UR_t = \beta(X_t - c'Z_{t-1}) + e_t = \beta(u_t + (d-c)'Z_{t-1}) + e_t = \beta u_t + \beta(d-c)'Z_{t-1} + e_t = \gamma'Z_{t-1} + s_t, \tag{5}$$

where $\gamma \equiv \beta(d-c)$ and $s_t \equiv \beta u_t + e_t$. Notice that s_t is orthogonal to Z_{t-1}: e_t is uncorrelated with Z_{t-1} by definition of the regression in the first line (and the fact that the parentheses do not impose a binding constraint) and u_t is uncorrelated with Z_{t-1} because it represents true unexpected earnings from the time-series regression in Eq. (4). It follows that γ is the slope when UR_t is regressed on Z_{t-1}. Thus, as long as β is not zero, the predictive slope on Z_{t-1} in a return regression differs from zero if and only if $d \neq c$. The implication is that everything we know about return predictability tests—the joint-hypothesis problem, the importance of controlling for cross-sectional correlation in the residuals, the impact of correlated omitted variables on the slopes, etc.—carries over directly to Mishkin tests.[3]

Since predictive regressions and Mishkin tests are nearly equivalent, the principal motivation for using a Mishkin test cannot be statistical but economic. Specifically, the Mishkin test provides a framework for *interpreting* predictability: rather than conclude simply that Z_{t-1} predicts returns ($\gamma \neq 0$), the Mishkin test interprets this predictability as coming from investors' misperceptions of the earnings process and quantifies how large this misperception must be ($c \neq d$). Indeed, this is precisely the way it has been used by Sloan (1996), Dechow et al. (2008), and others.

It is important to note, however, that even this motivation is questionable: The Mishkin test simply *assumes* that the predictive power of accruals (or whatever variable is being studied) comes from investors' misperceptions of the earnings process. It is entirely possible that investors forecast earnings correctly and, instead, misperceive the time-series process of, say, sales, dividends, cash flows, capital expenditures, or some other value-relevant variable. In other words, we could repeat the Mishkin test replacing earnings with any variable, Y_t, that is correlated with returns ($\beta \neq 0$):

$$UR_t = \beta(Y_t - c'Z_{t-1}) + e_t, \tag{6}$$

$$Y_t = d'Z_{t-1} + u_t, \tag{7}$$

---

[3] There is a subtle distinction here concerning the impact of correlated omitted variables: omitted variables may affect the slope on a particular predictive variable but do not affect the overall test of predictability (or market efficiency). For example, omitting B/M or past returns from a predictive regression (or a Mishkin test) may affect the slope on accruals but cannot make returns appear to be predictable when they are not. Kraft et al. (2007) discuss this issue in detail.

and it would be equally valid to interpret the predictive power of accruals as coming from investors' misperception of the time-series process of Y_t rather than of X_t. (In principle, we could include both X_t and Y_t in the return regression, but it is not possible to separately identify investors' misperceptions of the time-series processes for the two variables.)

Thus, using the Mishkin test to interpret the predictive power of Z_{t-1} relies heavily on a researcher's assumption of what drives predictability (e.g., biased forecasts of X_t) and does not actually provide evidence that misperceptions of X_t are truly the cause. Admittedly, earnings are a natural variable to consider, but, at the same time, a large empirical literature shows that contemporaneous earnings explain only a small fraction of annual stock returns. For example, RTW report an R^2 of 12% when annual returns are regressed on earnings changes from 1974 to 2009. Interpreting return predictability entirely through the lens of earnings forecasts makes the strong assumption that the remaining 88% of return variability plays no role.

As a concrete example, Sloan (1996) and Dechow et al. (2008) explore the predictive power of accruals for future earnings and returns. Both papers find that accruals predict stock returns, with similar predictive slopes on accruals, -0.28 for Sloan and -0.35 for Dechow et al. (these estimates of γ are calculated from the formula above, $\gamma \equiv \beta(d-c)$, using the values of β, d, and c provided in each paper). It turns out, however, that investors' misperceptions of the earnings process appear to be quite different in the two papers: in a regression of future earnings on lagged cash flow and accruals, Sloan estimates the true slope on accruals is 0.765 but the perceived slope is 0.911, implying that investors have an upward bias of 0.146; Dechow et al. estimate the true slope is 0.647 and the perceived slope is 0.938, implying an upward bias nearly twice as large, 0.291. (Accruals are defined differently in the papers, but Dechow et al. report that differences in the results are attributable primarily to different sample periods.) Thus, the papers find that accruals have roughly the same predictive power for returns but infer substantially different degrees of bias in investors' perceptions of the earnings process. These facts are reconciled by the different earnings response coefficients, β, reported by the two papers (1.894 by Sloan; 1.206 by Dechow et al.). While it is possible that a drop in β just happens to offset an increase in bias, so that $\gamma \equiv \beta(d-c)$ stays the same, an alternative explanation is that the accrual anomaly could be driven by investors' misperceptions of something other than earnings.

5. Research design: transaction costs and firm size

RTW's survey provides an extensive discussion of the important features of a well-designed empirical asset-pricing study. They suggest four key elements: (1) a credible null hypothesis, preferably with solid theoretical foundations; (2) rigorous statistical tests, with a special focus on the incremental predictive power of a variable relative to other anomalies; (3) a serious consideration of both risk and transaction costs; and (4) supplemental non-price tests—based, for example, on the properties of analyst forecast errors—as an independent way to check whether expectational biases can help explain any return predictability.

RTW's list provides a good blueprint for empirical studies, and it seems hard to disagree with any of their general suggestions. However, I believe it is useful to consider more closely their specific observations concerning transaction costs.

RTW rightly observe that 'an inference of market efficiency is quite different if the return magnitude is within the bounds of expected trading costs or not. If it is not, and can be reliably shown to be so, then this would provide more compelling evidence against efficient prices' (RTW, this issue). Thus, RTW's main concern here is whether an anomaly can be exploited profitably in practice.

I agree with RTW's basic argument, but my view of the role of transaction costs is somewhat different. Specifically, I would interpret an anomaly as providing evidence against market efficiency—assuming the pattern is truly anomalous—regardless of whether it can be exploited or not. Transaction costs are crucial for understanding how important the anomaly is and why it might not be arbitraged away by smart investors, but it seems too generous to say that prices are efficient simply because the anomaly cannot be profitably exploited. In addition, patterns like post-earnings-announcement drift and the accrual anomaly can help us to understand better how the market processes information even if the underlying predictability falls within the bounds of transaction costs. In short, anomalies are interesting even when they cannot be exploited by investors.

Let me give a specific example. The literature provides much evidence that small, less liquid stocks react with a partial delay to marketwide news (e.g., Lo and MacKinlay, 1990; Hou and Moskowitz, 2005). Regardless of whether the delay is large enough to trade on profitably, its existence suggests that (1) small stocks do not react efficiently to macroeconomic news, (2) investors who trade small stocks may be less sophisticated, and (3) substantial trading volume may be necessary for prices to fully reflect marketwide information. All of these inferences help us to understand the price-setting process and imply that small stocks can be at least temporarily mispriced—again, regardless of whether the pattern can be exploited by smart investors.

A separate concern is that it may be nearly impossible to accurately assess transaction costs. Even if we can estimate all of the direct and indirect trading costs discussed by RTW, we would still need to know an investor's entire portfolio—and other trading strategies—to assess the incremental cost of including an additional predictive variable in the investor's trading model. For example, studying the trading costs of a momentum strategy in isolation is much different than asking whether momentum might be profitably exploited by an investor who is already trading on value or accruals.

A related issue concerns the importance of firm size in cross-sectional asset-pricing tests. RTW suggest that it is useful to check whether an anomaly shows up in large stocks with low trading costs in order to test whether it is really exploitable. This concern is important, but I think that size plays a more fundamental role in asset-pricing tests, for at least two reasons. First, it is crucial to know whether an anomaly exists among large stocks in order to assess the anomaly's economic significance; predictability that is driven by small stocks is much less interesting because they make up only a small fraction of aggregate market value. Second, an interaction between size and mispricing can help us to understand what causes the anomaly in the first place, such as the behavioral biases of individuals vs. institutional traders or the role of analyst and media coverage (as well as transaction costs).

For these reasons, an important omission from RTW's blueprint is that it does not specifically advocate separate tests for small vs. large stocks. In my view, empirical studies should (1) report separate cross-sectional regressions for large stocks, not just pool all firms together; (2) focus on value-weighted returns in addition to (or instead of) equal-weighted returns; and (3) consider double-sorted portfolios, sorting firms by size and whatever predictive variable is being studied. Fama and French (2008) provide a useful classification scheme for this purpose: tiny stocks are those smaller than the NYSE 20th percentile, small stocks are those between the NYSE 20th and 50th percentiles, and large stocks are those bigger than the NYSE 50th percentile. (It is important to use NYSE breakpoints so that the percentiles are not dominated by the large number of tiny stocks.) At the end of 2009, the NYSE 20th, 50th, and 80th percentiles were \$416, \$1653, and \$6725 million, respectively (using all common stocks on CRSP with valid returns). These breakpoints do a good job of separating firms into the popular definitions of micro-cap vs. small-cap vs. mid-cap vs. large-cap stocks. The 2566 firms below the NYSE 20th percentile (57% of stocks) represent just 2% of overall market value, while the 977 firms above the NYSE median (22% of firms) represent 92% of overall market value.

6. RTW's empirical results

One of the nice features of RTW's survey is that it not only reviews the literature but also presents new empirical results for both the accrual anomaly and post-earnings-announcement drift. For accruals, RTW focus on the broad measure advanced by Fairfield et al. (2003), defined as the change in net operating assets (NOA) scaled by a firm's average total assets during the year (NOA equals non-cash total assets minus non-debt total liabilities). For earnings surprises, RTW focus on standardized unexpected earnings (SUE), defined as seasonally differenced earnings in the most recent quarter divided by the standard deviation of seasonally differenced earnings over the prior 12 quarters (earnings are measured before extraordinary items). The tests look at monthly returns from 1979 to 2008 for the 1000 largest stocks based on market value at the start of each month.

The main innovation of the empirical tests, relative to the literature, comes from the way RTW construct the trading strategies. The standard approach would be to form a long–short portfolio based on, say, the top and bottom deciles when stocks are sorted by accruals or SUE (see, e.g., Bernard and Thomas, 1990; Sloan, 1996). RTW instead construct a portfolio in a way that is designed to maximize its expected return net of risk and transaction costs. Specifically, they assume (for purposes of constructing the portfolio) that stocks' expected returns, μ, are linear in accruals or SUE and solve for the portfolio weights, w, that maximize a quadratic utility function:

$$w = \operatorname{argmax}_x x'\mu - \lambda x'Vx, \tag{8}$$

where V is an ex ante estimate of stocks' covariance matrix and λ is a risk-aversion parameter. For some tests, an additional term is included that penalizes high transaction-cost firms, where transaction costs are estimated using a proprietary model based on actual trade data. The optimal portfolio weights are solved numerically subject to the constraint that the weights sum to zero, the target standard deviation of the portfolio is 10% annually, and the maximum investment in any single stock is 5% positive or negative. The covariance matrix is estimated on a rolling basis, looking back, by breaking returns into a common factor component and an idiosyncratic component, imposing the assumption that the idiosyncratic component is uncorrelated across stocks.

RTW show that the resulting strategies perform quite well. Focusing on accruals, the optimized portfolio has a Fama–French three-factor alpha of 1.39% monthly with a *t*-statistic of 8.26 and a so-called appraisal ratio of 1.51 annually.[4] In comparison, an equal-weighted decile strategy has an alpha of 1.51% with a *t*-statistic of 2.99 and an appraisal ratio of 0.55, while a value-weighted decile strategy has an alpha of 0.65% with a *t*-statistic of 2.93 and an appraisal ratio of 0.54. RTW conclude that (1) the accrual anomaly is robust to ex ante controls for risk and transaction costs and (2) the risk- and transaction-cost-optimized strategy produces much stronger statistical evidence than the typical academic (decile-based) approach and a substantially better portfolio for trading purposes.

A few observations may be worthwhile. First, RTW's empirical tests are fully consistent with standard academic approaches: any asset-pricing test—whether it is based on cross-sectional regressions, decile sorts, or RTW's optimization

---

[4] The appraisal ratio is analogous to the Sharpe ratio but is defined as a portfolio's alpha divided by its residual risk (i.e., residual standard deviation in the factor regression). The appraisal ratio is (approximately) proportional to the *t*-statistic of the alpha, $t_\alpha \approx$ appraisal ratio $\times T^{1/2}$, just as the Sharpe ratio is (exactly) proportional to the *t*-statistic for average excess returns (see my discussion of Sharpe ratios in Section 2). I use that fact to infer the appraisal ratio from the *t*-statistics reported by RTW (the sample length, T, equals 358 months in RTW's tests). The ratio is annualized by multiplying it by the square root of 12.

procedure—boils down to asking whether some trading strategy's average return (or alpha) is significantly different from zero. For all of the discussion in RTW's paper about the importance of an 'ex ante consideration of risk' (RTW, this issue), ultimately the only way to evaluate their strategy is to study the portfolio's risk-return trade-off ex post.

At the same time, RTW's suggestion to consider ex ante measures of risk is fundamentally sound. It has two potential benefits. First, a key problem for asset-pricing tests is that they often have modest statistical power because returns are so noisy. The power of a test—which, as observed above, always boils down to asking whether some portfolio's average return is different from zero—is directly related to the Sharpe ratio of that portfolio: the bigger the true Sharpe ratio, the greater the power. (In statistical terms, the t-statistic has a t-distribution centered at zero under the null of no predictability but a non-central t-distribution under the alternative hypothesis, with a noncentrality parameter proportional to the portfolio's true Sharpe ratio.) Therefore, to the extent that we can maximize the ex ante risk-return trade-off of a portfolio, we can improve the power of an empirical test.

Consider a specific example. Suppose that returns are normally distributed and that a long–short strategy has a true Sharpe ratio of 0.30 annually, close to the market portfolio's historical Sharpe ratio. With 40 years of data, the t-statistic testing whether the strategy's expected return is different from zero is expected to be 1.90 ($0.30 \times 40^{1/2}$; see my earlier discussion of t-statistics vs. Sharpe ratios; the number I quote here is actually the noncentrality parameter of the t-distribution, which in small samples would be slightly different from the mean of the distribution). This test has a 48% probability of correctly rejecting the null at a 5% significance level (equal to the probability that the t-statistic is greater than 1.96 if it is drawn from a distribution with a mean of 1.90). If, by minimizing the risk of the trading strategy, we can increase the strategy's Sharpe ratio to 0.60, the expected t-statistic doubles to 3.79, and the power of the test jumps to 96% (the probability that the observed t-statistic is greater than 1.96 if it is drawn from a distribution with a mean of 3.79). This suggests that an increase in the risk-return trade-off of a portfolio can have a big impact on the power of the test.

A second important benefit of RTW's suggestion to optimize over risk is that we are almost always interested in how strong an anomaly is economically, not just statistically. One measure of that is a trading strategy's average return—which does not depend directly on risk—but a strategy's Sharpe ratio is probably a better metric. So, again, minimizing the ex ante risk of a portfolio could have a big impact on whether an anomaly appears to be economically interesting or not. And, as I argued earlier, the Sharpe ratio of a portfolio provides one way to distinguish risk from mispricing because risk stories have a hard time explaining the existence of extremely high Sharpe ratios.

In short, trying to minimize the risk of a trading strategy can potentially increase both the statistical and economic informativeness of an asset-pricing test. RTW's specific approach for doing so seems promising, more than doubling the Sharpe ratio and t-statistic of their accrual strategy compared with a decile-based strategy. However, it is important to note that RTW's strategy and the decile-based strategy differ in a couple of ways that have little to do with the fact that RTW's approach reflects an 'ex ante consideration of risk': (1) RTW's strategy potentially includes all 1000 stocks in the portfolio, while the decile-based strategy considers only the top and bottom 100 stocks; and (2) RTW's strategy partially adjusts for industry factors, since the risk model used for their optimization procedure includes industry fixed effects, while the decile-based strategy focuses on raw returns and raw accruals (not adjusted for industry factors). Thus, it is not clear how much of the improvement in RTW's strategy actually comes from their risk-optimization procedure.

Table 1 below provides tentative evidence on the importance of the two features mentioned above. Specifically, I replicate some features of RTW's tests but focus on basic portfolio strategies rather than their more elaborate optimization procedure. Following RTW, the sample includes the largest 1000 common stocks on CRSP ranked by market value at the beginning of each month (I exclude financial stocks but the results are similar if they are included). Accruals, based on annual Compustat data, equal the year-over-year change in NOA divided by average total assets during the year. I update accruals four months after the end of the fiscal year to ensure that accruals are publicly available (e.g., starting in May for a firm with a December fiscal year end; RTW update after three months and use quarterly Compustat data). For tests that require industry data, I define industries using Fama and French's 49 industry portfolios, available on Ken French's website at Dartmouth College; all stocks, regardless of market value, are included to get the industry mean in order to have a reasonable number of firms in each industry. I winsorize accruals monthly at the 1st and 99th percentiles, again including all stocks in the sample when I calculate percentiles.

The table shows the performance of six trading strategies, three using raw accruals and returns and three using industry-adjusted accruals and returns (subtracting the industry mean). All three strategies in each panel represent zero-investment portfolios with one dollar long and one dollar short. The first two are based on value-weighted (VW) and equal-weighted (EW) deciles (decile 1 minus decile 10). The third is a 'linear-weight' portfolio with weights that are proportional to accruals, i.e., $w_i = p(\Delta NOA_i - \mu_{\Delta NOA})$, where w_i is stock i's portfolio weight, $\mu_{\Delta NOA}$ is the cross-sectional mean of accruals, and p is a constant that scales the weights up or down to have one dollar long and one dollar short (p is negative, implying that any stock with below-average accruals gets a positive weight and any stock with above-average accruals gets a negative weight). The various strategies are chosen to explore the benefit of including all stocks in the portfolio (the linear-weight strategy), not just the extreme deciles, and to explore whether industry adjustments have much impact on the results.

The table shows several interesting results. First, the EW decile strategies and the linear-weight strategies have similar performance: the EW strategies have slightly higher returns and alphas, but the linear-weight strategies have lower standard deviations and slightly more significant alphas (as mentioned earlier, an alpha's t-statistic is proportional to the strategy's appraisal ratio, a measure similar to the Sharpe ratio except that it is defined using risk-adjusted returns).

Table 1

This table reports the average monthly return (Avg, %), standard deviation (Std, %), Sharpe ratio, CAPM alpha (%), and Fama–French (FF) three-factor alpha (%) for six accrual trading strategies from 1979 to 2008. Strategies in the top panel are based on raw accruals and returns while strategies in the bottom panel are based on industry-adjusted accruals and returns (subtracting the industry mean, determined using Fama and French's 49 industry classification). Each panel shows three strategies, the first based on value-weighed (VW) deciles, the second based on equal-weighted (EW) deciles, and the third based on a linear-weight portfolio. The decile strategies buy the bottom decile and short the top decile of stocks sorted by lagged accruals. The linear-weight strategy invests in stocks in proportion to their accruals (the weights are scaled and shifted to have one dollar long and one dollar short). The sample includes the largest 1000 nonfinancial common stocks on CRSP. Accruals are measured as the annual change in net operating assets scaled by average total assets during the year (updated four months after the end of the fiscal year).

Strategy	Raw returns			CAPM alpha		FF alpha	
	Avg.	Std	Sharpe [a]	a	$t(a)$	a	$t(a)$
Strategies based on raw accruals and returns							
VW deciles	0.62	3.46	0.62	0.76	4.36	0.52	3.17
EW deciles	0.85	3.17	0.93	1.01	6.57	0.74	5.46
Linear weights	0.74	2.92	0.88	0.91	6.86	0.66	6.12
Strategies based on industry-adjusted accruals and returns							
VW deciles	0.55	3.74	0.51	0.73	4.08	0.61	3.39
EW deciles	0.75	2.70	0.97	0.89	6.97	0.73	6.04
Linear weights	0.62	2.02	1.06	0.72	7.68	0.59	6.75

[a] Annualized.

This result suggests that focusing on just the extreme accrual deciles does not result in a substantially lower risk-return trade-off than a strategy that uses all stocks. Second, even within the universe of the largest 1000 stocks, the EW strategies perform substantially better than the VW strategies—as measured by average returns, Sharpe ratios, or alphas—suggesting that the accrual anomaly is weaker among the very largest stocks. Third, my industry adjustment has some impact on the results but improves performance only slightly (as measured by the Sharpe ratios or *t*-statistics).

The performance of my decile strategies is notably different from the results reported by RTW, despite the seemingly small differences in the way our portfolios are constructed (e.g., the use of annual vs. quarterly Compustat data). That finding suggests that, while the basic accrual anomaly is quite robust, inferences about its exact strength may depend on modest differences in methodology. Perhaps most interesting, *t*-statistics for the Fama–French alphas of my EW decile strategy (5.46 without the industry adjustment; 6.04 with) are closer to the *t*-statistic for RTW's risk-optimized strategy (8.26) than to the *t*-statistic for their EW decile strategy (2.99). It seems plausible, then, that part of the reason RTW's risk-optimized strategy performs so much better than their EW decile strategy is simply that their EW portfolio exhibits surprisingly high variability.

7. Conclusions

Empirical capital markets research remains an active and important area of study. Over the past decade, the literature has produced, by RTW's count, 165 new papers that have significantly improved our understanding of how accounting numbers relate to firm value and how investors use—and potentially misinterpret—accounting numbers. These issues have important implications for investors, creditors, firm managers, and policy makers.

This article provides a critical review of five issues that are central to the literature on accounting anomalies and fundamental analysis:

(1) *Risk vs. mispricing:* Interpreting stock price behavior is one of the most challenging aspects of empirical research. Without a full model of how prices should be set in an efficient market, it is difficult to know whether return predictability reflects risk or mispricing. The most rigorous way to disentangle the two is to ask whether predictability is consistent with any reasonable risk-based model, the cleanest rejection coming if a pattern represents an arbitrage opportunity.

(2) *Implied cost of capital:* The high volatility of stock prices means that sample estimates of expected returns are imprecise. This problem has motivated new ways of estimating expected returns, one of which is to infer a firm's cost of capital from observed stock prices and forecasts of future dividends and earnings. This approach makes sense in principle, but I am skeptical of how useful it is in practice because of the difficulty of forecasting dividends and earnings, the challenge of empirically testing measures of the implied cost of capital, and the fact that the implied cost of capital represents at best the long-run internal rate of return from holding a stock, not the short-run expected return or the appropriate discount rate for a firm's investment decisions.

(3) *Mishkin tests:* A popular way to interpret accounting anomalies is to compare the true time-series properties of earnings (estimated using observed data) to the properties perceived by investors (inferred from stock returns). However, these tests *assume* that an anomaly reflects investors' misperceptions of the earnings process and do not actually provide evidence that it does so.

(4) *Transaction costs and firm size:* If returns are predictable, it is interesting to know whether investors can take advantage of the predictability after accounting for trading costs. But an anomaly may still be interesting even if it cannot be profitably exploited. A more important issue, in my view, is to test whether an anomaly exists only among smaller firms—which make a large fraction of firms but only a small fraction of market value—or whether it is also present among the economically more important population of large stocks.

(5) *Risk optimization:* RTW argue that researchers, like practitioners, should consider ex ante measures of risk when they construct stock portfolios. Their suggestion is fundamentally sound, not because ex ante measures of risk are more important than the ex post riskiness of the portfolio, but because risk optimization can improve a test's power and our ability to discriminate between risk and mispricing. The practical benefits of risk optimization are an important topic for future research.

RTW provide a separate review of these, and many other, issues. I have tried to complement their study and to provide useful suggestions for future research.

Acknowledgement

References

Abel, Andrew, Mishkin, Frederic, 1983. An integrated view of tests of rationality, market efficiency, and the short-run neutrality of monetary policy. Journal of Monetary Economics 11, 3–24.

Acharya, Viral, Pedersen, Lasse, 2005. Asset pricing with liquidity risk. Journal of Financial Economics 77, 375–410.

Ait-Sahalia, Yacine, Parker, Jonathan, Yogo, Motohiro, 2004. Luxury goods and the equity premium. Journal of Finance 59, 2959–3004.

Bansal, Ravi, Dittmar, Robert, Lundblad, Christian, 2005. Consumption, dividends, and the cross section of equity returns. Journal of Finance 60, 1639–1672.

Berk, Jonathan, Green, Richard, Naik, Vasant, 1999. Optimal investment, growth options, and security returns. Journal of Finance 54, 1553–1607.

Bernard, Victor, Thomas, Jacob, 1990. Evidence that stock prices do not fully reflect the implications of current earnings for future earnings. Journal of Accounting and Economics 13, 305–340.

Botosan, Christine, Plumlee, Marlene, 2005. Assessing alternative proxies for the expected risk premium. The Accounting Review 80, 21–53.

Bradshaw, Mark, Richardson, Scott, Sloan, Richard, 2001. Do analysts and auditors use information in accruals? Journal of Accounting Research 39, 45–74.

Brennan, Michael, Wang, Ashley, Xia, Yihong, 2004. Estimation and test of a simple model of intertemporal asset pricing. Journal of Finance 59, 1743–1775.

Campbell, John, Cochrane, John, 1999. By force of habit: a consumption-based explanation of aggregate stock market behavior. Journal of Political Economy 107, 205–251.

Campbell, John, Cochrane, John, 2000. Explaining the poor performance of consumption-based asset pricing models. Journal of Finance 55, 2863–2878.

Campbell, John, Vuolteenaho, Tuomo, 2004. Bad beta, good beta. American Economic Review 94, 1249–1275.

Carlson, Murray, Fisher, Adlai, Giammarino, Ron, 2004. Corporate investment and asset price dynamics: implications for the cross-section of returns. Journal of Finance 59, 2577–2603.

Carlson, Murray, Fisher, Adlai, Giammarino, Ron, 2006. Corporate investment and asset price dynamics: implications for SEO event studies and long-run performance. Journal of Finance 61, 1009–1034.

Cochrane, John, 1996. A cross-sectional test of an investment-based asset pricing model. Journal of Political Economy 104, 572–621.

Cohen, Randolph, Polk, Christopher, Vuolteenaho, Tuomo, 2009. The price is (almost) right. Journal of Finance 64, 2739–2782.

Dechow, Patricia, Richardson, Scott, Sloan, Richard, 2008. The persistence and pricing of the cash component of earnings. Journal of Accounting Research 46, 537–566.

Dechow, Patricia, Sloan, Richard, 1997. Returns to contrarian investment strategies: tests of naïve expectations hypotheses. Journal of Financial Economics 43, 3–27.

Duffee, Gregory, 2005. Time variation in the covariance between stock returns and consumption growth. Journal of Finance 60, 1673–1712.

Easton, Peter, 2004. PE ratios, PEG ratios, and estimating the implied expected rate of return on equity capital. The Accounting Review 79, 73–95.

Easton, Peter, Monahan, Steven, 2005. An evaluation of accounting-based measures of expected returns. The Accounting Review 80, 501–538.

Easton, Peter, Monahan, Steven, 2010. Evaluating accounting-based measures of expected returns: Easton and Monahan and Botosan and Plumlee Redux. Working paper (University of Notre Dame).

Easton, Peter, Sommers, Gregory, 2007. Effect of analysts' optimism on estimates of the expected rate of return implied by earnings forecasts. Journal of Accounting Research 45, 983–1015.

Fairfield, Patricia, Whisenant, J. Scott, Yohn, Teri, 2003. Accrued earnings and growth: implications for future profitability and market mispricing. The Accounting Review 78, 353–371.

Fama, Eugene, 1970. Efficient capital markets: a review of theory and empirical work. Journal of Finance 25, 383–417.

Fama, Eugene, 1976. Foundations of Finance. Basic Books, New York, NY.

Fama, Eugene, French, Kenneth, 1992. The cross-section of expected stock returns. Journal of Finance 47, 427–465.

Fama, Eugene, French, Kenneth, 1993. Common risk factors in the returns on stocks and bonds. Journal of Financial Economics 33, 3–56.

Fama, Eugene, French, Kenneth, 2008. Dissecting anomalies. Journal of Finance 63, 1653–1678.

Gebhardt, William, Lee, Charles, Swaminathan, Bhaskaran, 2001. Toward an implied cost of capital. Journal of Accounting Research 39, 135–176.

Gomes, Joao, Kogan, Leonid, Zhang, Lu, 2003. Equilibrium cross section of returns. Journal of Political Economy 111, 693–732.

Hail, Luzi, Leuz, Christian, 2009. Cost of capital effects and changes in growth expectations around US cross-listings. Journal of Financial Economics 93, 428–454.

Hansen, Lars, Heaton, John, Li, Nan, 2008. Consumption strikes back? Measuring long-run risk. Journal of Political Economy 116, 260–302.

Hansen, Lars, Jagannathan, Ravi, 1991. Implications of security market data for models of dynamic economies. Journal of Political Economy 99, 225–262.

Heaton, John, Lucas, Deborah, 2000. Asset pricing and portfolio choice: the role of entrepreneurial risk. Journal of Finance 55, 1163–1198.

Hou, Kewei, Moskowitz, Tobias, 2005. Market frictions, price delay, and the cross-section of expected returns. Review of Financial Studies 18, 981–1020.

Jacobs, Kris, Wang, Kevin, 2004. Idiosyncratic consumption risk and the cross section of asset returns. Journal of Finance 59, 2211–2252.

Jagannathan, Ravi, Wang, Zhenyu, 1996. The conditional CAPM and the cross section of stock returns. Journal of Finance 51, 3–53.

Jegadeesh, Narasimhan, Titman, Sheridan, 1993. Returns to buying winners and selling losers: implications for stock market efficiency. Journal of Finance 48, 65–91.

Jegadeesh, Narasimhan, Titman, Sheridan, 2001. Profitability of momentum strategies: an evaluation of alternative explanations. Journal of Finance 56, 699–720.

Kraft, Arthur, Leone, Andrew, Wasley, Charles, 2007. Regression-based tests of the market pricing of accounting numbers: the Mishkin test and ordinary least squares. Journal of Accounting Research 45, 1081–1114.

La Porta, Rafael, 1996. Expectations and the cross-section of stock returns. Journal of Finance 51, 1715–1742.

La Porta, Rafael, Lakonishok, Josef, Shleifer, Andrei, Vishny, Robert, 1997. Good news for value stocks: further evidence on market efficiency. Journal of Finance 52, 859–874.

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

Lettau, Martin, Ludvigson, Sydney, 2001. Resurrecting the (C)CAPM: a cross-sectional test when risk premia are time-varying. Journal of Political Economy 109, 1238–1287.

Li, Erica, Livdan, Dmitry, Zhang, Lu, 2009. Anomalies. Review of Financial Studies 22, 4301–4334.

Li, Qing, Vassalou, Maria, Xing, Yuhang, 2006. Sector investment growth rates and the cross section of equity returns. Journal of Business 89, 1637–1665.

Lo, Andrew, MacKinlay, A.Craig, 1990. When are contrarian profits due to stock market overreaction? Review of Financial Studies 3, 175–205

Lustig, Hanno, Van Nieuwerburgh, Stijn, 2004. Housing collateral, consumption insurance, and risk premia. Journal of Finance 60, 1167–1221.

Mehra, Rajnish, Prescott, Edward, 1985. The equity premium: a puzzle. Journal of Monetary Economics 15, 145–161.

Merton, Robert, 1973. An intertemporal asset pricing model. Econometrica 41, 867–887.

Mishkin, Frederic, 1983. A Rational Expectations Approach to Macroeconometrics: Testing Policy Ineffectiveness and Efficient Markets Models. National Bureau of Economic Research Monograph (University of Chicago Press, Chicago, IL).

Novy-Marx, Robert, 2010. Operating leverage. Review of Finance 14, 1–32.

Ohlson, James, 2005. On accounting-based valuation formulae. Review of Accounting Studies 10, 322–347.

Ohlson, James, Juettner-Nauroth, Beate, 2005. Expected EPS and EPS growth as determinants of value. Review of Accounting Studies 10, 349–365.

Parker, Jonathan, Julliard, Christian, 2005. Consumption risk and the cross section of expected returns. Journal of Political Economy 113, 185–222.

Pastor, Lubos, Stambaugh, Robert, 2003. Liquidity risk and expected stock returns. Journal of Political Economy 111, 642–685.

Petkova, Ralitsa, 2006. Do the Fama–French factors proxy for innovations in predictive variables? Journal of Finance 61, 581–612

Richardson, Scott, Sloan, Richard, Soliman, Mark, Tuna, Irem, 2005. Accrual reliability, earnings persistence, and stock prices. Journal of Accounting and Economics 39, 437–485.

Richardson, Scott, Sloan, Richard, Soliman, Mark, Tuna, Irem, 2006. The implications of firm growth and accounting distortions for accruals and profitability. The Accounting Review 81, 713–743.

Richardson, Scott, Tuna, Irem, Wysocki, Peter, this issue. Accounting anomalies and fundamental analysis: a review of recent research advances. Journal of Accounting and Economics, doi:10.1016/j.jacceco.2010.09.008.

Santos, Tano, Veronesi, Pietro, 2006. Labor income and predictable stock returns. Review of Financial Studies 19, 1–44.

Sloan, Richard, 1996. Do stock prices fully reflect information in accruals and cash flows about future earnings? The Accounting Review 71, 289–315.

Stein, Jeremy, 1996. Rational capital budgeting in an irrational world. Journal of Business 69, 429–455.

Vassalou, Maria, 2003. News related to future GDP growth as a risk factor in equity returns. Journal of Financial Economics 68, 47–73.

Wu, Jin, Zhang, Lu, Zhang, X. Frank, 2010. The q-theory approach to understanding the accrual anomaly. Journal of Accounting Research 48, 177–223.

Yogo, Motohiro, 2006. A consumption-based explanation of expected stock returns. Journal of Finance 61, 539–580.

Zhang, Lu, 2005. The value premium. Journal of Finance 60, 67–103.