

On the Predictability of Stock Returns: Theory and Evidence

by

Jonathan W. Lewellen

Submitted in Partial Fulfillment

of the

Requirements for the Degree

Doctor of Philosophy

Supervised by

Professor Jay Shanken

William E. Simon Graduate School of Business Administration

University of Rochester
Rochester, New York

2000

Curriculum Vitae

The author was born in West Lafayette, Indiana on May 1, 1972. He studied economics and finance at Indiana University from 1990 to 1994, and graduated with a Bachelor of Science degree in 1994. He came to the University of Rochester in the Fall of 1994 and began graduate studies in applied economics and finance at the Simon Graduate School of Business Administration. He received financial support from the Simon School from 1994 to 1999 and received an Olin Fellowship from 1997 to 1998. He pursued his research, specializing in the pricing of financial assets, under the direction of professors G. William Schwert, Jay Shanken, and Jerold Warner. He earned a Master of Science in applied economics in 1997.

Acknowledgements

I gratefully acknowledge the advice and support of my dissertation committee, Bill Schwert, Jay Shanken, and Jerry Warner. Their teaching and encouragement are ultimately responsible for this dissertation. I owe a special debt to Jay Shanken, whose guidance has stimulated my research in countless ways and whose friendship is yet more valuable. I have also learned much from my colleagues at the Simon School, especially Katharina Bibus, Andreas Ginschel, Jarrad Harford, Christoph Hinkelmann, Michelle Lowry, Susan Shu, Denis Suvorov, and Peter Wysocki. This paper has benefited from the comments of Greg Bauer, Kent Daniel, Ken French, S.P. Kothari, John Long, René Stulz, and Jiang Wang. I thank the Simon Graduate School of Business Administration and the Olin Foundation for financial support.

Abstract

Empirical studies find that stock returns are predictable both cross-sectionally and over time. Broadly speaking, this dissertation investigates whether the empirical patterns in stock returns are consistent with an efficient capital market. The paper consists of two essays. In the first essay, I investigate the ability of firms' book-to-market ratios to predict returns, which has been documented extensively in cross-sectional tests. To help understand the source of this predictability, I examine the time-series relations among expected return, risk, and book-to-market. Consistent with rational pricing, book-to-market captures significant time-variation in risk, but provides no incremental information about expected returns. In the second essay, I explore the effects of estimation risk, or investor uncertainty about the parameters of the cashflow process, on the behavior of prices and returns. I show that, with estimation risk, the observable properties of prices and returns can differ significantly from the properties perceived by rational investors. As a consequence, estimation risk can generate return predictability in ways that resemble irrational pricing. Simulation evidence suggests the effects of estimation risk can be economically significant.

Table of Contents

Chapter 1	Introduction	1
Chapter 2	The Time-Series Relations among Expected Return, Risk, and Book-to-Market	3
2.1.	Distinguishing between Characteristics and Risk.....	8
2.2.	Data and Descriptive Statistics	13
2.3.	The Predictability of Portfolio Returns	18
2.4.	Expected Returns, Characteristics, and Risk: Empirical Results	25
2.5.	Summary and Conclusions	40
Chapter 3	Estimation Risk, Market Efficiency, and the Predictability of Returns	43
3.1.	The Model.....	49
3.2.	Capital Market Equilibrium	53
3.3.	The Time-Series Properties of Prices and Returns	58
3.4.	The Cross-Section of Expected Returns	67
3.5.	Informative Priors, Steady State, and Simulations	73
3.6.	Summary and Conclusions	89
References.....		93
Appendix A.....		98
Appendix B		104

List of Tables

Chapter 2

Table 2.1	Summary statistics for industry, size, and book-to-market portfolios	15
Table 2.2	Summary statistics for factors	18
Table 2.3	Predictability of industry returns	21
Table 2.4	Predictability of size and book-to-market portfolio returns	24
Table 2.5	Unconditional three-factor regressions: Industry portfolios	27
Table 2.6	Conditional three-factor regressions: Industry portfolios	29
Table 2.7	Unconditional three-factor regressions: Size and book-to-market portfolios	32
Table 2.8	Conditional three-factor regressions: Size and book-to-market portfolios	35
Table 2.9	Three-factor regressions with industry-neutral HML: Industry portfolios	39

Chapter 3

Table 3.1	Predictability in steady state	87
-----------	--------------------------------------	----

List of Figures

Chapter 3

Figure 3.1 Equilibrium price of the risky asset..... 61

On the Predictability of Stock Returns: Theory and Evidence

Chapter 1 Introduction

Over the past 20 years, we have accumulated much evidence that stock returns are predictable. At the aggregate level, Fama and Schwert (1977), Keim and Stambaugh (1986), Fama and French (1989), and Kothari and Shanken (1997) show that interest rates, the yield spread between low- and high-grade debt, aggregate dividend yield, and aggregate book-to-market predict time-variation in expected returns. Further, LeRoy and Porter (1981) and Shiller (1981) argue that the volatility of stock prices is too high to be explained by a model with constant discount rates, providing indirect evidence that expected returns change over time. At the firm level, Fama and French (1992) conclude that size and book-to-market together explain much of the cross-sectional variation in average returns. Jegadeesh and Titman (1993) also show that past returns contain additional information about expected returns. In sum, there seems little doubt that expected stock returns vary both cross-sectionally and over time.¹

The interpretation of predictability, however, is more contentious. The empirical patterns in returns are potentially consistent with either market efficiency or irrational mispricing. In general terms, market efficiency implies that prices ‘fully reflect all available information.’ To formalize this idea for empirical testing, Fama (1976) distinguishes between the probability distribution of returns perceived by ‘the market,’ based on whatever information investors view as relevant, and the true distribution of returns conditional on all information. The market is said to be informationally efficient if these distributions are the same. As an obvious consequence, market efficiency implies that investors correctly anticipate any cross-sectional or time-variation in true expected returns. While Fama’s definition ignores potentially important issues like heterogeneous beliefs, it provides a useful framework for thinking about a broad set

¹ Clearly, this list of empirical papers and predictive variables is not meant to be exhaustive, and a considerable amount of subsequent research extends, confirms, and critiques these findings. See Fama (1991) for a more complete survey of the evidence.

of asset-pricing questions.

This paper contains two essays which, broadly speaking, attempt to understand whether the empirical results are consistent with an efficient capital market. In the first essay, I investigate the ability of ‘book-to-market’ ratios to predict returns. An extensive literature shows that the ratio of a firm’s book value to market value of equity – in short, book-to-market – explains significant cross-sectional variation in expected returns. The intuition is that book value in the numerator controls for the size of the firm (the size of expected cashflows), while market value in the denominator captures information about discount rates. Both efficient-market and mispricing stories have been offered to explain the evidence. In this paper, I examine the time-series relations among expected return, risk, and book-to-market to help understand the source of the predictability. As discussed in Chapter 2, the time-series analysis can help distinguish between the rational- and irrational-pricing stories.

In the second essay, I investigate the impact of ‘estimation risk’ on the behavior of asset prices. In the finance literature, estimation risk refers to investor uncertainty about the parameters of the return or cashflow process. In other words, estimation risk exists whenever investors do not have perfect information about some important feature of the economy. Although it represents purely subjective uncertainty, estimation risk can have important consequences for asset pricing because it affects investment decisions. In Chapter 3, I present a simple model of capital market equilibrium, and explore the consequences of estimation risk for return predictability and tests of market efficiency. I also present simulation evidence to give an indication of the economic significance of the results. Again, the fundamental goal is to understand whether estimation risk might help explain the time-series and cross-sectional evidence described above.

On the Predictability of Stock Returns: Theory and Evidence

Chapter 2

The time-series relations among expected return, risk, and book-to-market

Empirical research consistently finds a positive cross-sectional relation between average stock returns and the ratio of a firm's book equity to market equity (B/M). Stattman (1980) and Rosenberg, Reid, and Lanstein (1985) document the association between expected returns and B/M, which remains significant after controlling for beta, size, and other firm characteristics (Fama and French, 1992). The explanatory power of B/M does not appear to be driven entirely by data snooping or survival biases; it is found in stock markets outside the United States (Chan, Hamao, and Lakonishok, 1991; Haugen and Baker, 1996) and in samples drawn from sources other than Compustat (Davis, 1994). As a whole, the evidence provides considerable support for the cross-sectional explanatory power of B/M.

At least two explanations have been offered for the empirical evidence. According to asset-pricing theory, B/M must proxy for a risk factor in returns. The significance of B/M in competition with beta contradicts the capital asset pricing model (CAPM) of Sharpe (1964), Lintner (1965), and Black (1972), or more precisely, the mean-variance efficiency of the market proxy. However, the evidence might be consistent with the intertemporal models of Merton (1973) and Breeden (1979). In these models, the market return does not completely capture the relevant risk in the economy, and additional factors are required to explain expected returns. If a multifactor model accurately describes stock returns, and B/M is cross-sectionally correlated with the factor loadings, then the premium on B/M simply reflects compensation for risk.

A positive relation between B/M and risk is expected for several reasons. Chan and Chen (1991) and Fama and French (1993) suggest that a distinct 'distress factor' explains common variation in stock returns. Poorly performing, or distressed, firms are likely to have high B/M. These firms are especially sensitive to economic conditions, and their returns might be driven

by many of the same macroeconomic factors (such as variation over time in bankruptcy costs and access to credit markets). In addition, following the arguments of Ball (1978) and Berk (1995), B/M might proxy for risk because of the inverse relation between market value and discount rates. Holding book value constant in the numerator, a firm's B/M ratio increases as expected return, and consequently risk, increases.

Alternatively, B/M might provide information about security mispricing. The mispricing view takes the perspective of a contrarian investor. A firm with poor stock price performance tends to be underpriced and have a low market value relative to book value. As a result, high B/M predicts high future returns as the underpricing is eliminated. Lakonishok, Shleifer, and Vishny (1994) offer a rationale for the association between past performance and mispricing. They argue that investors naively extrapolate past growth when evaluating a firm's prospects. For example, investors tend to be overly pessimistic about a firm which has had low or negative earnings. On average, future earnings exceed the market's expectation, and the stock does abnormally well. Thus, the mispricing argument says that B/M captures biases in investor expectations.

Fama and French (1993) provide evidence of a relation between B/M and risk. Using the time-series approach of Black, Jensen, and Scholes (1972), they examine a multifactor model consisting of market, size, and book-to-market factors, where the size and book-to-market factors are stock portfolios constructed to mimic underlying risk factors in returns. If the model explains cross-sectional variation in average returns, the intercepts will be zero when excess returns are regressed on the three factors. Fama and French find, as predicted by the risk-based view, that the model does a good job explaining average returns for portfolios sorted by size, B/M, earnings-price ratios, and other characteristics. Further, they document a strong association between a stock's B/M ratio and its loading on the book-to-market factor.

More recently, Daniel and Titman (1997) argue in favor of a characteristics-based model, consistent with the mispricing view. They suggest that the three-factor model does not *directly*

explain average returns. Instead, the model appears to explain average returns only because the factor loadings are correlated with firms' characteristics (size and B/M). To disentangle the explanatory power of the factor loadings from that of the characteristics, Daniel and Titman construct test portfolios by sorting stocks first on B/M ratios and then on factor loadings. This sorting procedure creates independent variation in the two variables. Consistent with the mispricing story, Daniel and Titman find a stronger relation between expected returns and B/M than between expected returns and factor loadings. Daniel and Titman conclude that firm characteristics, in particular B/M, and not covariances determine expected stock returns.

In this essay, I provide further evidence on the risk- and characteristics-based stories. In contrast to Fama and French (1993) and Daniel and Titman (1997), I focus on the time-series relations among expected return, risk, and B/M. Specifically, I ask whether a portfolio's B/M ratio predicts time-variation in its expected return, and test whether changes in expected return can be explained by changes in risk. Recently, Kothari and Shanken (1997) and Pontiff and Schall (1998) find that B/M forecasts stock returns at the aggregate level, but the predictive ability of B/M for individual stocks or portfolios has not been explored.

The time-series analysis is a natural alternative to cross-sectional regressions. An attractive feature of the time-series regressions is that they focus on changes in expected returns, not on average returns. The mispricing story suggests that a stock's expected return will vary over time with B/M, but it says little about average returns if mispricing is temporary. Cross-sectional regressions, however, can pick up a relation between average returns and B/M. The time-series regressions also highlight the interaction between B/M and risk, as measured by time-variation in market betas and the loadings on the Fama and French (1993) size and book-to-market factors. Further, I can directly test whether the three-factor model explains time-varying expected returns better than the characteristics-based model. These results should help distinguish between the risk and mispricing stories.

The empirical tests initially examine B/M's predictive ability without attempting to control

for changes in risk. I find that a portfolio's B/M ratio tracks economically and statistically significant variation in its expected return. An increase in B/M equal to twice its time-series standard deviation forecasts a 4.6% (annualized) increase in expected return for the typical industry portfolio, 8.2% for the typical size portfolio, and 9.3% for the typical book-to-market portfolio. The average coefficient on B/M across all portfolios, 0.99, is approximately double the cross-sectional slope, 0.50, found by Fama and French (1992, p. 439). B/M explains, however, only a small fraction of portfolio returns, generally less than 2% of total volatility.

Return predictability indicates that either risk or mispricing changes over time. Of course, we cannot distinguish between these explanations without some model of risk. Following Daniel and Titman (1997), I examine B/M's explanatory power in competition with the Fama and French (1993) three-factor model. The multifactor regressions employ the conditional asset-pricing methodology of Shanken (1990), which allows both expected returns and factor loadings to vary over time with B/M. In these regressions, time-variation in the intercepts measures the predictive ability of B/M that cannot be explained by changes in risk. The mispricing view suggests that the intercepts will be positively related to B/M; the risk-based view implies that changes in the factor loadings will eliminate B/M's explanatory power, assuming the Fama and French factors are adequate proxies for priced risk in the economy.

Empirically, the factors absorb much of the volatility of portfolio returns, which permits relatively powerful tests of the competing stories. I find that B/M explains significant time-variation in risk, but does not provide incremental information about expected return. In general, the loadings on the size and book-to-market factors vary positively with a portfolio's B/M ratio, and statistical tests strongly reject the hypothesis of constant risk. The results for market betas are more difficult to characterize: across different portfolios, B/M predicts both significant increases and significant decreases in beta. Overall, B/M contains substantial information about the riskiness of stock portfolios.

In contrast, the intercepts of the three-factor model do not vary over time with B/M. For

the industry portfolios, the average coefficient on B/M (that is, variation in the intercept) has the opposite sign predicted by the overreaction hypothesis and is not significantly different from zero. Across the 13 portfolios, eight coefficients are negative and none is significantly positive at conventional levels. The results are similar for size and book-to-market portfolios: the average coefficients are indistinguishable from zero, and roughly half are negative. Importantly, the inferences from the multifactor regressions are not driven by low power. For all three sets of portfolios, statistical tests can reject economically large coefficients on B/M. In short, the three-factor model measures risk sufficiently well to explain time-variation in expected returns.²

As an aside, I find that the book-to-market factor, HML, explains common variation in returns that is unrelated to its industry composition. Daniel and Titman (1997) argue that HML does not proxy for a distinct risk factor, but explains return covariation only because similar types of firms become mispriced at the same time. For example, a bank with high B/M will covary positively with HML simply because the factor is weighted towards underpriced financial firms. The time-series regressions provide evidence to the contrary. As an alternative to HML, I estimate the regressions with an ‘industry-neutral’ book-to-market factor. This factor is constructed by sorting stocks on their industry-adjusted B/M ratios, defined as the firm’s B/M minus the industry average, so the factor should never be weighted towards particular industries. The results using the industry-neutral factor are similar to those with HML. Thus, HML’s explanatory power does not appear to be driven by industry factors in returns.

The remainder of the essay is organized as follows. Section 2.1 introduces the time-series regressions. Section 2.2 describes the data to be used in the empirical tests. Section 2.3 estimates the simple relation between expected returns and B/M, and Section 2.4 tests whether

² I also replicate the empirical tests using size in place of B/M, with similar results. There is some evidence that size and expected returns are negatively related in time series. In conditional three-factor regressions, size captures significant time-variation in risk, but does not contain additional information about expected returns. Details are available on request. I thank Ken French for suggesting these tests.

the predictive ability of B/M can be explained by changes in risk, as measured by the Fama and French (1993) three-factor model. Section 2.5 summarizes the evidence and concludes.

2.1. Distinguishing between characteristics and risk

Book-to-market explains cross-sectional variation in average returns after controlling for beta. Fama and French (1993) provide evidence that B/M relates to common risk factors in returns. In contrast, Daniel and Titman (1997) argue that the Fama and French factors appear to be priced only because the loadings are correlated with firm characteristics, like B/M. This section introduces the time-series methodology used in the current paper and discusses, more generally, asset-pricing tests of the risk and mispricing stories.

2.1.1. Time-series methodology

The empirical tests initially examine the simple relation between expected returns and B/M. The explanations that have been offered for the cross-sectional evidence also suggest that expected returns will vary over time with B/M. According to the risk-based view, B/M should capture information about changes in risk, and consequently, expected return. The mispricing view says that B/M is related to biases in investor expectations, and will contain information about under- and overpricing. Thus, both explanations predict a positive slope coefficient in the regression

$$R_i(t) = \gamma_0 + \gamma_1 B/M_i(t-1) + e_i(t), \quad (2.1)$$

where R_i is the portfolio's excess return and B/M_i is its lagged book-to-market ratio. Note that eq. (2.1) specifies a separate time-series regression for each portfolio, with no constraint on the coefficients across different portfolios. The regressions focus only on the time-series relation between expected returns and B/M, and do not pick up any cross-sectional relation.

Eq. (2.1) makes no attempt to understand the source of time-varying expected returns. According to traditional asset-pricing theory, a positive slope in eq. (2.1) must be driven by an

association between B/M and risk. It follows that the predictive power of B/M should be eliminated if the regressions control adequately for changes in risk. The characteristics-based story, on the other hand, suggests that B/M will capture information about expected returns that is unrelated to risk. To help distinguish between the two explanations, I examine the predictive power of B/M in competition with the Fama and French (1993) three-factor model.

The multifactor regressions employ the conditional time-series methodology of Shanken (1990). Roughly speaking, these regressions combine the three-factor model with the simple regressions above. Fama and French estimate the unconditional model

$$R_i(t) = a_i + b_i R_M(t) + s_i \text{SMB}(t) + h_i \text{HML}(t) + e_i(t), \quad (2.2)$$

where R_M is the excess market return, SMB (small minus big) is the size factor, and HML (high minus low) is the book-to-market factor. Unconditional, here, refers to the implicit assumption that the coefficients of the model are constant over time. If this assumption is not satisfied, the estimates from eq. (2.2) can be misleading. The unconditional intercepts and factor loadings could be close to zero, but might vary considerably over time.

The conditional regressions allow both expected returns and factor loadings to vary with B/M. Suppose, for simplicity, that the coefficients of the three-factor model are linearly related to the firm's B/M ratio, or

$$\begin{aligned} a_{it} &= a_{i0} + a_{i1} B/M_i(t-1), & b_{it} &= b_{i0} + b_{i1} B/M_i(t-1), \\ s_{it} &= s_{i0} + s_{i1} B/M_i(t-1), & h_{it} &= h_{i0} + h_{i1} B/M_i(t-1). \end{aligned} \quad (2.3)$$

Substituting these equations into the unconditional regression yields a conditional version of the three-factor model:

$$\begin{aligned} R_i &= a_{i0} + a_{i1} B/M_i + (b_{i0} + b_{i1} B/M_i) * R_M + \\ &\quad (s_{i0} + s_{i1} B/M_i) * \text{SMB} + (h_{i0} + h_{i1} B/M_i) * \text{HML} + e_i, \end{aligned} \quad (2.4)$$

where the time subscripts have been dropped to reduce clutter. Multiplying the factors through gives the regression equation for each portfolio. Thus, the conditional regressions contain not

only an intercept and the three factors, but also four interactive terms with the portfolio's lagged B/M.³

Basically, eq. (2.4) breaks the predictive power of B/M into risk and non-risk components. The coefficient a_1 , the interactive term with the intercept, measures the predictive ability of B/M that is incremental to its association with risk in the three-factor model. A non-zero coefficient says that changes in the factor loadings, captured by the coefficients b_{i1} , s_{i1} , and h_{i1} , do not fully explain the time-series relation between B/M and expected return. Thus, rational asset-pricing theory predicts that a_1 will be zero for all stocks, assuming that the factors are adequate proxies for priced risk. The mispricing, or characteristics-based, view implies that B/M will forecast returns after controlling for risk and, consequently, a_{i1} should be positive.

2.1.2. Discussion

The conditional regressions directly test whether the three-factor model or the characteristic-based model better explains changes in expected returns. To interpret the regressions as a test of rational pricing, we must assume, of course, that the Fama and French factors capture priced risk in the economy. This assumption could be violated in two important ways (see Roll, 1977). First, an equilibrium multifactor model might describe stock returns, but the Fama and French factors are not adequate proxies for the unknown risks. In this case, B/M can predict time-variation in expected returns missed by the three-factor model if it relates to the true factor loadings. Fortunately, this problem will not be a concern for the current paper because the three-factor model will, in fact, explain the predictability associated with B/M.

Unfortunately, the assumption can also be violated in the opposite way: mispricing might explain deviations from the CAPM, but the size and book-to-market factors happen to absorb

³ Similar regressions appear in previous studies. Fama and French (1997) estimate regressions in which only the factor loadings on HML vary with B/M. He et al. (1996) estimate a model in the spirit of eq. (2.4), but they constrain the intercepts and book-to-market coefficients to be the same across portfolios. Given previous cross-sectional evidence, the B/M coefficient will be non-zero in the absence of time-varying expected returns.

the predictive power of B/M. This possibility is a concern particularly because the factors are empirically motivated. Daniel and Titman (1997), for example, argue that the construction of HML, which is designed to mimic an underlying risk factor in returns related to B/M, could induce ‘spurious’ correlation between a portfolio’s B/M ratio and its factor loading. HML is weighted, by design, towards firms with high B/M. If similar types of firms become mispriced at the same time, then we should expect that a firm will covary more strongly with HML when its B/M is high. As a result, apparent changes in risk might help explain B/M’s predictive ability even under the mispricing story.

In defense of the time-series regressions, it seems unlikely that changes in the factor loadings would completely absorb mispricing associated with B/M. More importantly, Daniel and Titman’s argument cannot fully account for the relation between B/M and risk. The argument suggests that the loadings on HML will tend to vary with B/M, but it does not say anything about the loadings on the market and size factors. We will see below, however, that B/M captures significant time variation in market betas and the loadings on SMB. Further, I provide evidence in Section 2.4 that the time-series relation between B/M and the factor loadings on HML is not driven by changes in the industry composition of the factor. I estimate the conditional regressions with an ‘industry neutral’ factor, which prevents HML from becoming weighted towards particular industries. When this factor is used in place of HML, we will continue to see a strong time-series relation between B/M and the factor loadings.

Finally, it is useful to note that many industries have large *unconditional* factor loadings on HML, which suggests that HML does not simply capture mispricing in returns. Intuitively, Daniel and Titman’s argument suggests that a given stock will sometimes vary positively and sometimes negatively with HML. Depending on the type of firms that are currently under- and overpriced, HML will be related to constantly changing micro- and macroeconomic factors. For example, HML will be sensitive to interest rate and inflation risk when it is weighted towards underpriced financial firms, but will be negatively related to these risks when financial

firms are overpriced. Corresponding to the changes in HML, a stock will tend to covary positively with HML when similar firms are underpriced, but negatively when similar firms are overpriced. Over time, however, a firm's average factor loading on HML should be close to zero under the mispricing story, unless firms are persistently under- and overpriced (which seems unreasonable).

This intuition can be formalized. Suppose that temporary overreaction explains deviations from the CAPM, and that HML, because of its construction, absorbs this mispricing (ignore the size factor for simplicity). To be more specific, assume that the proxy for the market portfolio, M , is not mean-variance efficient conditional on firms' B/M ratios. However, HML is constructed to explain the deviations from the CAPM, and R_M and HML together span the conditional tangency portfolio. Appendix A proves that, in the time-series regression

$$R_i(t) = a_i + b_i R_M(t) + h_i \text{HML}(t) + e_i(t), \quad (2.5)$$

the unconditional factor loading on HML, h_i , will equal zero if assets are correctly priced on average over time.⁴ This result reflects the idea that temporary mispricing should not explain unconditional deviations from the CAPM. As noted above, however, many industries have large unconditional loadings on both SMB and HML, which therefore suggests that the factors do not simply capture mispricing in returns.

In summary, the multifactor regressions test whether the three-factor model or the characteristic-based model explains time-variation in expected returns. The interpretation of the regressions, like the results for any asset-pricing test, is limited by our need to use a proxy for the unobservable model. Nevertheless, the regressions should help us understand whether the risk or mispricing story is a better description of asset prices.

⁴ The result also requires that time-variation in b_i and h_i is uncorrelated with the factors' expected returns. This assumption seems reasonable since I am interested in the factor loadings changing over time with firm-specific variables, like B/M, not with macroeconomic variables (the appendix provides a numerical example). It is also consistent with the empirical evidence presented in Section 2.4.

2.2. Data and descriptive statistics

The empirical analysis focuses on industry portfolios. These portfolios should exhibit cross-sectional variation in expected returns and risk, so the tests can examine a diverse group of portfolios. Industry portfolios are believed a priori to provide variation in expected returns and factor loadings, while sorting by other criteria is often motivated by previous empirical evidence. Hence, industry portfolios are less susceptible to the data-snooping issues discussed by Lo and MacKinlay (1990).

As a robustness check, I also examine portfolios sorted by size and B/M. In cross-sectional studies, different sets of portfolios often produce vastly different estimates of risk premia. Of course, the time-series regressions in this paper might also be sensitive to the way portfolios are formed. Size portfolios have the advantage that they control for changes in market value, which has been shown to be associated with risk and expected returns, yet should be relatively stable over time. The book-to-market portfolios allow us to examine how the expected returns and risk of distressed, or high-B/M, firms change over time.

The portfolios are formed monthly from May 1964 through December 1994, for a time series of 368 observations. The industry and size portfolios consist of all NYSE, Amex, and Nasdaq stocks on the Center for Research in Security Prices (CRSP) tapes, while the book-to-market portfolios consist of the subset of stocks with Compustat data. Stocks are sorted into 13 industry portfolios based on two-digit Standard Industrial Classification (SIC) codes as reported by CRSP. For the most part, the industries consist of consecutive two-digit codes, although some exceptions were made when deemed appropriate.⁵ The size portfolios are formed based on the market value of equity in the previous month, with breakpoints determined by NYSE deciles. To reduce the fraction of market value in any single portfolio, the largest two portfolios are further divided based on the 85th and 95th percentiles of NYSE stocks, for a total of 12 portfolios. Finally, the book-to-market portfolios are formed based on the ratio of book equity

⁵ Details available on request.

in the previous fiscal year to market equity in the previous month. Again, the breakpoints for these portfolios are determined by NYSE deciles. The lowest and highest deciles are further divided using the 5th and 95th percentiles of NYSE stocks, for a total of 12 portfolios.

For all three sets of portfolios, value-weighted returns are calculated using all stocks with CRSP data, and value-weighted B/M ratios are calculated from the subset of stocks with Compustat data.⁶ To ensure that the explanatory power of B/M is predictive, I do not assume that book data become known until five months after the end of the fiscal year. Also, to reduce the effect of potential selection biases in the way Compustat adds firms to the database (see the discussion by Kothari, Shanken, and Sloan, 1995), a firm must have three years of data before it is included in any calculation requiring book data. The time-series regressions use excess returns, calculated as returns minus the one-month T-bill rate, and the natural logarithm of B/M.

Table 2.1 reports summary statistics for the portfolios. The average monthly returns for the industry portfolios range from 0.83% for utilities and telecommunications firms to 1.28% for the service industry (which includes entertainment, recreation, and services), for an annualized spread of 6.1%. Coincidentally, these industries also have the lowest (3.67%) and highest (6.78%) standard deviations, respectively. The size and book-to-market portfolios also exhibit wide variation in average returns and volatility. Average returns for the size portfolios vary from 0.80% for the largest stocks to 1.24% for the smallest stocks, and the standard deviations of returns decrease monotonically with size, from 6.68% to 4.17%. Average returns for the book-to-market portfolios range from 0.76% for the second decile through 1.46% for the stocks with the highest B/M. Interestingly, the standard deviation of returns are U-shaped; they decrease monotonically with B/M until the sixth decile, which has a standard deviation of 4.42%, and increase thereafter, to 6.86% for portfolio 10b.

The statistics for B/M, like those for returns, reveal considerable cross-sectional

⁶ The stocks included in the calculation of B/M are a subset of those included in the calculation of returns, and we can interpret the estimate of B/M as a proxy for the entire portfolio. The inferences in this paper are unchanged when portfolio returns are based only on those stocks with Compustat data.

Table 2.1

Summary statistics for industry, size, and book-to-market portfolios

Each month from May 1964 through December 1994, value-weighted portfolios are formed monthly from all NYSE, Amex, and Nasdaq stocks on CRSP. Firms must also have Compustat data for the book-to-market portfolios. Book-to-market (B/M) is calculated as the ratio of book equity in the previous fiscal year to market equity in the previous month for all stocks with Compustat data. The industry portfolios are based on two-digit SIC codes. The size portfolios are based on the market value of equity in the previous month, with breakpoints determined by NYSE deciles; portfolios 9 and 10 are further divided using the 85 and 95 percentiles of NYSE stocks. The book-to-market portfolios are based on B/M in the previous month, with breakpoints determined by NYSE deciles; portfolios 1 and 10 are further divided using the 5 and 95 percentiles of NYSE stocks.

Portfolio	Return (%)		Book-to-market				Number of firms	
	Mean	Std. dev.	Mean	Std. dev.	Autocorr.	Adj. R ^{2a}	May 1964	Dec. 1994
<i>Panel A: Industry portfolios</i>								
Nat. resources	0.84	5.69	0.57	0.14	0.97	0.28	90	360
Construction	0.86	5.45	0.78	0.26	0.99	0.69	237	409
Food, tobacco	1.21	4.57	0.51	0.18	0.99	0.36	106	134
Consumer products	1.05	6.02	0.75	0.36	0.99	0.81	108	257
Logging, paper	0.98	5.35	0.54	0.15	0.98	0.54	74	190
Chemicals	0.96	4.78	0.40	0.13	0.99	0.28	102	392
Petroleum	1.08	5.23	0.74	0.20	0.98	0.38	30	32
Machinery, equipment	0.88	5.35	0.42	0.13	0.99	0.40	290	1,222
Transportation	0.87	5.39	0.82	0.27	0.98	0.71	162	260
Utilities, telecom.	0.83	3.67	0.77	0.25	0.99	0.57	122	384
Trade	1.04	5.65	0.49	0.18	0.98	0.55	167	785
Financial	0.95	4.75	0.75	0.19	0.97	0.68	117	1,747
Services and other	1.28	6.78	0.47	0.20	0.98	0.55	61	981
<i>Panel B: Size portfolios</i>								
Smallest	1.24	6.68	1.03	0.41	0.98	0.88	409	3,338
2	1.16	6.17	0.90	0.32	0.97	0.79	175	932
3	1.15	6.08	0.85	0.30	0.97	0.77	154	634
4	1.16	5.91	0.84	0.31	0.97	0.80	149	463

Table 2.1. Continued.

Portfolio	Return (%)		Book-to-market				Number of firms	
	Mean	Std. dev.	Mean	Std. dev.	Autocorr.	Adj. R ² ^a	May 1964	Dec. 1994
5	1.21	5.70	0.76	0.26	0.96	0.80	136	426
6	1.22	5.41	0.72	0.23	0.97	0.88	130	333
7	1.07	5.19	0.68	0.19	0.97	0.88	132	310
8	1.09	5.13	0.67	0.19	0.98	0.86	130	274
9a	1.02	4.89	0.67	0.18	0.97	0.78	62	122
9b	0.97	4.71	0.66	0.20	0.97	0.78	64	111
10a	0.88	4.50	0.62	0.17	0.98	0.70	63	109
Largest	0.80	4.17	0.51	0.15	0.99	0.41	62	109
<i>Panel C: Book-to-market portfolios</i>								
Lowest	0.98	5.69	0.15	0.05	0.97	0.59	21	559
1b	0.83	5.15	0.24	0.07	0.97	0.64	19	328
2	0.76	5.04	0.34	0.10	0.97	0.82	40	493
3	0.79	4.78	0.46	0.14	0.98	0.93	39	470
4	0.83	4.63	0.57	0.18	0.98	0.94	39	469
5	0.82	4.47	0.67	0.21	0.98	0.94	42	492
6	0.98	4.42	0.78	0.24	0.98	0.94	40	467
7	1.17	4.54	0.89	0.28	0.98	0.94	42	461
8	1.25	4.72	1.04	0.32	0.98	0.95	43	460
9	1.43	5.20	1.28	0.40	0.98	0.96	43	574
10a	1.46	6.12	1.65	0.53	0.97	0.93	24	378
Highest	1.46	6.86	2.66	0.95	0.96	0.84	25	482

^a Adjusted R² from regressing the portfolio's B/M ratio on the value-weighted B/M ratio of all stocks that meet both CRSP and Compustat data requirements.

differences in portfolio characteristics. Average B/M doubles from 0.40 for chemical firms to 0.82 for the transportation industry. A similar spread is shown for size portfolios, with B/M ranging from 0.51 for the largest stocks to 1.03 for the smallest stocks. The book-to-market portfolios, of course, have the greatest cross-sectional variation, with average B/M ranging from 0.15 for the low-B/M portfolio to 2.66 for the high-B/M portfolio. The standard deviations over time are also reasonably high, reflecting the volatility of stock returns. The time-series standard deviation of B/M is, on average, 0.20 for the industries, 0.24 for the size portfolios, and 0.29 for the book-to-market portfolios. Variation in B/M will be necessary for the time-series regressions to have power distinguishing between the competing hypotheses.

Table 2.2 reports summary statistics for the Fama and French (1993) factors, which are described fully in Appendix A. The market factor, R_M , is the excess return on the CRSP value-weighted index, and the size and book-to-market factors, SMB and HML, are zero-investment portfolios designed to mimic underlying risk factors in returns. The average monthly return of R_M is 0.39%, of SMB is 0.30%, and of HML is 0.38%. The risk premium for each factor is measured by its mean return, so these averages imply positive compensation for bearing factor risk. As noted by Fama and French, the procedure used to construct SMB and HML appears to successfully control each factor for the influence of the other, as demonstrated by the low correlation between the factors, equal to -0.06. Also, SMB is positively correlated with R_M (correlation of 0.36), while HML is negatively correlated with R_M (-0.35). Thus, the returns on the size and B/M factors are not independent of the market return, reflecting the fact that their construction did not control for differences in the betas of the underlying stocks.

The CAPM and most empirical studies examine the relation between simple-regression market betas and expected returns. To enhance comparison with cross-sectional studies, I use size and B/M factors that are orthogonal to R_M . These factors, SMBO and HMLO, are constructed by adding the intercepts to the residuals when SMB and HML are regressed on a constant and the excess market return. From regression analysis (e.g., Johnston, 1984, p. 238),

Table 2.2
Summary statistics for factors

The factors are calculated monthly from May 1964 through December 1994. R_M is the return on the CRSP value-weighted index minus the one-month T-bill rate. SMB is the return on a portfolio of small stocks minus the return on a portfolio of big stocks. HML is the return on portfolio of high-B/M stocks minus the return on a portfolio of low-B/M stocks. SMBO and HMLO are orthogonalized versions of SMB and HML, constructed by adding the intercepts to the residuals in regressions of SMB and HML on a constant and R_M . All returns are reported in percent.

Factor	Mean	Std. dev.	Autocorr.	Correlation				
				R_M	SMB	HML	SMBO	HMLO
R_M	0.39	4.45	0.06	1.00	0.36	-0.35	0.00	0.00
SMB	0.30	2.91	0.19		1.00	-0.06	0.93	0.07
HML	0.38	3.00	0.14			1.00	0.07	0.94
SMBO	0.21	2.71	0.06				1.00	0.07
HMLO	0.47	2.81	0.14					1.00

the coefficients in the three-factor model will be unaffected by the change in variables, except that market betas will now be the simple-regression betas of the CAPM. Table 2.2 shows that the average return on the book-to-market factor increases from 0.38% to 0.47%, but the return on the size factor decreases from 0.30% to 0.21%. The correlation between the size and book-to-market factors, 0.07, remains close to zero.

2.3. The predictability of portfolio returns

This section investigates the simple time-series relation between expected returns and B/M. The simple regressions help evaluate the economic importance of B/M, without regard to changes in risk or mispricing, and provide a convenient benchmark for the conditional three-factor model. In addition, the analysis complements recent studies which find that B/M forecasts aggregate stock returns (Kothari and Shanken, 1997; Pontiff and Schall, 1998).

As discussed above, the risk and mispricing views both suggest that B/M will predict portfolio returns. For each portfolio, I estimate the time-series regression

$$R_i(t) = \gamma_0 + \gamma_1 B/M_i(t-1) + e_i(t), \quad (2.6)$$

where R_i is the portfolio's excess return and B/M_i is the natural log of its lagged book-to-market ratio. The slope coefficient in this regression is expected to be positive.

Several complications arise in estimating eq. (2.6). First, the appropriate definition of B/M is unclear. Cross-sectional studies suggest that a portfolio's B/M relative to other firms could be important. Thus, $B/M_i(t-1)$ might be defined as either the portfolio's actual B/M ratio or its B/M ratio minus an aggregate index. The latter varies primarily with market-adjusted stock returns, and would be a better measure if common variation in B/M is unrelated to mispricing.⁷ Asset-pricing theory provides little guidance. The conclusions in this paper are not sensitive to the definition of B/M, and for simplicity I report only results for raw B/M. Also, to ease the interpretation of the results, B/M is measured as deviations from its time-series mean for the remainder of the paper. As a consequence, when B/M_i equals zero in the regressions, B/M is actually at its long-run average for the portfolio.

Second, Stambaugh (1999) shows that contemporaneous correlation between returns and B/M will bias upward the slope coefficient in eq. (2.6). Suppose that B/M follows the AR(1) process

$$B/M_i(t) = c_i + p_i B/M_i(t-1) + u_i(t). \quad (2.7)$$

The bias in the estimate of γ_{i1} is approximately

$$E[\hat{\gamma}_{i1} - \gamma_{i1}] \approx [\text{cov}(e_i, u_i) / \text{var}(u_i)] \cdot [-(1+3p_i) / T], \quad (2.8)$$

where T is the length of the time series. The residuals in eqs. (2.6) and (2.7), e_i and u_i , are negatively related because a positive stock return decreases the portfolio's B/M. Also, Table 2.1 shows that B/M is highly persistent over time, with autocorrelations ranging from 0.96 to 0.99 at the first lag. Together, the correlation between e_i and u_i and the persistence in B/M

⁷ Kothari and Shanken (1997) and Pontiff and Schall (1998) show that aggregate B/M predicts market returns during the period 1926 through 1992, which could reflect aggregate mispricing. Their results for the period 1963 through 1992 are much weaker. For the current paper, preliminary tests indicate that aggregate B/M has little power to forecast the market, size, and book-to-market factors.

impart a strong upward bias in the estimate of γ_{it} . In a related context, for market returns regressed on aggregate B/M, Kothari and Shanken (1997) bootstrap the distribution of the slope and find that Stambaugh's formula is empirically valid. The tests below adjust for this bias.

2.3.1. *Industry portfolios*

Table 2.3 reports results for the industry portfolios. The evidence provides some support for a positive association between expected returns and lagged B/M, but the high volatility of stock returns reduces the power of the tests. The bias-adjusted slopes range from -0.53 for food and tobacco firms to 1.75 for the natural resources industry, and 10 of the 13 coefficients are greater than zero. The average estimate is positive, 0.58, although it is only about one standard error, 0.62, from zero (the standard error reflects cross-sectional correlation in the estimates). Stronger evidence of predictive ability is provided by the χ^2 test of the slope coefficients. This test rejects at the 5% level the hypothesis that B/M does not capture any variation in expected returns.

The average coefficient, 0.58, is similar to the cross-sectional slope, 0.50, estimated by Fama and French (1992). Economically, the average coefficient is reasonably large. Consider, for example, the effect that a change in B/M equal to two standard deviations would have on expected returns. For the average industry portfolio, the time-series standard deviation of B/M is 0.33. An increase in B/M twice this large maps into a 0.38% change (0.66×0.58) in expected return for the typical portfolio, or 4.67% annually. On the other hand, the predictive power of B/M is low as measured by the adjusted R^2 s. Lagged B/M explains at most 1% of the total variation in portfolio returns. This result is consistent with previous studies at the market level, which generally find that pre-determined variables explain only a small fraction of monthly returns (e.g., Fama and French, 1989).

In addition to the ordinary least squares (OLS) estimates just described, Table 2.3 reports seemingly unrelated regression (SUR) estimates of the equations. OLS treats the regression for

Table 2.3
Predictability of industry returns

$$R_i(t) = \gamma_{i0} + \gamma_{i1} B/M_i(t-1) + e_i(t)$$

The industry portfolios are described in Table 2.1. R_i is the portfolio's monthly excess return (in percent) and B/M_i is the natural log of the portfolio's book-to-market ratio at the end of the previous month. The table reports both ordinary least squares (OLS) and seemingly unrelated regression (SUR) estimates of the slope coefficients. The OLS bias-adjusted slopes correct for small-sample biases using eq. 2.8 in the text. The bias correction for the SURs, as well as the covariance matrix of the bias-adjusted estimates, is obtained from bootstrap simulations.

Portfolio	OLS				SUR			
	γ_{i1}	Bias-adj γ_{i1}	Std. err.	Adj. R^2	γ_{i1}	Std. err.	Bias-adj γ_{i1}	Std. err.
Nat. resources	2.56*	1.75	1.19	0.01	0.67	0.69	0.07	0.77
Construction	1.10	0.15	0.83	0.00	0.13	0.29	-0.30	0.44
Food, tobacco	0.35	-0.53	0.64	0.00	0.50	0.31	0.17	0.43
Consumer products	0.88	-0.06	0.72	0.00	0.71*	0.31	0.33	0.45
Logging, paper	2.07	1.20	1.07	0.01	0.20	0.43	-0.30	0.55
Chemicals	1.01	0.09	0.78	0.00	0.06	0.35	-0.41	0.51
Petroleum	2.33*	1.46	1.00	0.01	1.56*	0.67	0.96	0.80
Mach., equipment	0.77	-0.19	0.80	0.00	-0.25	0.31	-0.69	0.50
Transportation	1.19	0.37	0.80	0.00	0.41	0.35	-0.02	0.47
Utilities, telecom.	1.06	0.40	0.57	0.01	1.11	0.36	0.77	0.52
Trade	1.69	0.97	0.88	0.01	0.70	0.37	0.35	0.45
Financial	1.93*	1.18	0.96	0.01	1.14*	0.42	0.74	0.47
Services, other	1.69	0.79	0.93	0.01	0.86*	0.36	0.51	0.46
Average	1.43*	0.58			0.60*		0.17	
(std. err.)	(0.62)	(0.62)			(0.20)		(0.23)	
χ^2 ^a	18.44	26.42*			22.48*		9.83	
(p-value)	(0.142)	(0.015)			(0.048)		(0.707)	

^a $\chi^2 = c' \Sigma^{-1} c$, where c is the vector of coefficient estimates and Σ is the estimate of the covariance matrix of c . Under the null that all coefficients are zero, this statistic is asymptotically distributed as χ^2 (d.f. 13).

* Denotes coefficients that are greater than two standard errors from zero or χ^2 statistics with a p-value less than 0.050.

each portfolio separately, and ignores interactions among the equations. The residuals across portfolios are correlated, however, because industries' excess returns are driven by many of the same macroeconomic factors. SUR uses this information to estimate the system of equations more efficiently (Zellner, 1962). Although SUR requires an estimate of the residual covariance matrix, the efficiency gain is likely to be large because (1) the error terms are highly correlated across portfolios (see Greene, 1993, p. 489), and (2) the dimension of the covariance matrix (13×13) is small relative to the length of the time series (368 months). Indeed, Table 2.3 shows that the average standard deviation of the SUR slopes is 0.40, compared with 0.86 for OLS. While the standard deviations are estimated with error, the large decrease suggests that SUR is substantially more efficient.

It was noted above that OLS slope estimates are biased upward. I am not aware of any research that explores the bias in SUR estimates, and there is little reason to believe that it is identical to that of OLS. Without an analytical estimate, I rely on bootstrap simulations to assess the sampling distribution of the SUR slopes. The simulation procedure, described in Appendix A, randomly generates time series of returns and B/M, imposing the restriction that expected returns and B/M are unrelated. Since the true coefficient in the simulation equals zero, the mean of the distribution represents the bias in SUR estimates. Further, the standard deviation of the distribution provides an estimate of the SUR standard error.⁸

Table 2.3 shows that the bias-adjusted SUR estimates tend to be smaller than their OLS counterparts. The coefficients range from -0.69 for the machinery and equipment industry to 0.96 for petroleum firms, and eight of the 13 estimates are positive. The average coefficient on B/M, 0.17, is positive, although it is under one standard error, 0.23, from zero. In addition, the χ^2 statistic cannot reject the hypothesis that all slope coefficients are zero. The simulations

⁸ I also simulate the distribution of the OLS slope estimates and find that the analytical estimate of the bias is reasonably accurate. The average bias from the simulations is 0.92 compared with 0.85 from eq. (2.8). The standard errors from the simulation, however, tend to be larger than the OLS estimates. For example, the standard deviation of the average coefficient is 0.76, compared with the OLS standard error of 0.62.

indicate that the average bias in the SUR estimates, 0.43, is about half the bias in the OLS regressions, 0.85. The magnitude remains significant, however, and the average SUR coefficient decreases by two-thirds, from 0.60 to 0.17, after correcting for bias.

In sum, the evidence in Table 2.3 is consistent with a positive relation between B/M and expected returns, but B/M explains, at most, a small fraction of returns. After adjusting for bias in the regressions, only the χ^2 statistic for the OLS slope coefficients is significant at conventional levels. We will see below that the power of the tests is much greater in the conditional three-factor regressions, because the factors absorb much of the volatility of returns. In addition, the size and book-to-market portfolios reveal a considerably stronger relation between B/M and future returns.

As a final observation, it is useful to keep in mind that the regressions cannot reject economically meaningful coefficients on B/M. A typical confidence interval around the average estimate, for either OLS or SUR, would include reasonably large coefficients. Moreover, low explanatory power does not imply that B/M is necessarily unimportant. For example, Kandel and Stambaugh (1996) show that predictive variables with low explanatory power can have a large impact on asset allocation decisions. I suspect a similar result would hold at the portfolio level: the optimal portfolio held by a risk-averse, Bayesian investor is probably sensitive to predictive variables which have low statistical significance.

2.3.2. Size and book-to-market portfolios

Table 2.4 shows results for the size and book-to-market portfolios. For simplicity, I report only the SUR estimates, along with the bias-adjusted estimates, since the evidence above indicates that SUR increases the precision of the slope estimates. The table shows that B/M predicts statistically reliable variation in returns for both the size and book-to-market portfolios. After correcting for bias, four coefficients for the size portfolios and nine coefficients for the book-to-market portfolios are more than two standard errors above zero. All 12 estimates are

Table 2.4

Predictability of returns: Size and book-to-market portfolios

$$R_i(t) = \gamma_{i0} + \gamma_{i1} B/M_i(t-1) + e_i(t)$$

The size and book-to-market portfolios are described in Table 2.1. R_i is the portfolio's monthly excess return (in percent) and B/M_i is the natural log of the portfolio's book-to-market ratio at the end of the previous month. The table reports seemingly unrelated regression (SUR) estimates of the slope coefficients, together with bias-adjusted slope estimates. The bias correction for the SURs, as well as the covariance matrix of the bias-adjusted estimates, is obtained from bootstrap simulations.

Size portfolios					Book-to-market portfolios				
Portfolio	SUR γ_{i1}	Std. err.	Bias-adj γ_{i1}	Std. err.	Portfolio	SUR γ_{i1}	Std. err.	Bias-adj γ_{i1}	Std. err.
Smallest	0.07	0.33	0.00	0.28	Lowest	-0.22	0.57	-0.45	0.54
2	0.32	0.23	0.31*	0.15	1b	-0.22	0.55	-0.46	0.50
3	0.30	0.21	0.28*	0.13	2	0.57	0.55	0.31	0.45
4	0.31	0.19	0.29	0.15	3	1.21*	0.51	0.96*	0.44
5	0.18	0.20	0.15	0.15	4	1.32*	0.50	1.09*	0.42
6	0.19	0.20	0.16	0.15	5	1.32*	0.51	1.06*	0.45
7	0.38	0.22	0.34*	0.17	6	1.19*	0.52	0.93*	0.45
8	0.45*	0.22	0.40*	0.18	7	1.80*	0.53	1.51*	0.46
9a	0.46	0.26	0.42	0.21	8	1.85*	0.56	1.53*	0.49
9b	0.44	0.24	0.38	0.21	9	2.09*	0.60	1.77*	0.49
10a	0.54	0.28	0.45	0.24	10a	2.27*	0.74	1.88*	0.61
Largest	0.18	0.33	0.03	0.39	Highest	2.52*	0.79	2.14*	0.68
Average	0.32*		0.27*		Average	1.31*		1.02*	
(std. err.)	(0.16)		(0.08)		(std. err.)	(0.46)		(0.29)	
χ^2 ^a	10.84		20.71		χ^2 ^a	25.12*		28.21*	
(p-value)	(0.543)		(0.055)		(p-value)	(0.014)		(0.005)	

^a $\chi^2 = c' \Sigma^{-1} c$, where c is the vector of coefficient estimates and Σ is the estimate of the covariance matrix of c . Under the null that all coefficients are zero, this statistic is asymptotically distributed as χ^2 (d.f. 12).

* Denotes coefficients that are greater than two standard errors from zero or χ^2 statistics with a p-value less than 0.050.

positive for the size portfolios, and the average coefficient, 0.27, is greater than three standard errors from zero. Similarly, ten of the 12 coefficients for the book-to-market portfolios are positive, and the average coefficient, 1.02, is more than three standard errors above zero. The estimates generally increase from the low-B/M deciles to the high-B/M deciles.

Interestingly, the bias in the regressions is significantly smaller for the size and book-to-market portfolios than for the industry portfolios. The bootstrap estimate of the bias is 0.05 for the size portfolios and 0.29 for the book-to-market portfolios, compared with 0.43 for the industries (see Table 2.3). Also, the standard errors from the simulated distribution are less than the actual SUR estimates, while the opposite is true for industry portfolios. From the bootstrap distribution, the standard error of the average coefficient is only 0.08 for the size portfolios and 0.29 for the book-to-market portfolios.

Economically, the individual estimates and the average coefficient are quite large for the book-to-market portfolios. A two-standard-deviation increase in B/M for the typical portfolio predicts a 0.61% monthly increase in expected return, or 7.6% annually. The implied change in expected return is greater than 11% annually for the five portfolios with the highest B/M. The conclusions from the OLS regressions (not reported) are qualitatively similar, but the estimates are less precise. The average bias-adjusted OLS slope is 1.13 (standard error of 0.82) for the size portfolios and 1.30 (standard error of 0.77) for the book-to-market portfolios. The strong relation between expected returns and B/M documented in Table 2.4 should provide a challenging test of the three-factor model.

2.4. Expected returns, characteristics, and risk: Empirical results

The evidence above indicates the B/M predicts significant time-variation in expected returns. In this section, I examine the explanatory power of B/M in competition with the Fama and French (1993) three-factor model. As discussed above, the conditional regressions directly test whether the three-factor model or the characteristic-based model better explains changes in

expected returns over time.

Fama and French estimate the unconditional model

$$R_i(t) = a_i + b_i R_M(t) + s_i \text{SMBO}(t) + h_i \text{HMLO}(t) + e_i(t), \quad (2.9)$$

where SMB and HML have been replaced here by the orthogonalized factors SMBO and HMLO (see Section 2.2). The conditional version of the three-factor model allows the intercepts and factor loadings to vary linearly with lagged B/M. Repeating eq. (2.4), the conditional model is specified as:

$$R_i = a_{i0} + a_{i1} B/M_i + (b_{i0} + b_{i1} B/M_i) R_M + (s_{i0} + s_{i1} B/M_i) \text{SMBO} + (h_{i0} + h_{i1} B/M_i) \text{HMLO} + e_i, \quad (2.10)$$

where B/M is lagged one month relative to returns and time subscripts have been dropped for simplicity. Multiplying the factors through gives the equation to be estimated for each portfolio. The B/M interactive term with the intercept, a_{i1} , is analogous to the slope coefficient in the simple regressions above, except that the multifactor regressions control for changes in risk. Consequently, a_{i1} measures the predictive ability of B/M that cannot be explained by the Fama and French three-factor model.

2.4.1. Industry portfolios

Before continuing to the conditional model, Table 2.5 reports unconditional three-factor regressions for the industry portfolios.⁹ Consistent with the results of Fama and French (1997), the size and book-to-market factors explain significant co-movement in industry returns not captured by the market. For both SMBO and HMLO, ten of the 13 coefficients deviate from zero by more than two standard errors. In fact, nine coefficients on the size factor and eight coefficients on the book-to-market factor are greater than four standard errors from zero. If the loadings change over time and are uncorrelated with the factors, the unconditional estimates can

⁹ For these regressions, OLS and SUR are identical because the regressors are the same for all portfolios (Greene, 1993, p. 488).

Table 2.5

Unconditional three-factor regressions: Industry portfolios

$$R_i(t) = a_i + b_i R_M(t) + s_i \text{SMBO}(t) + h_i \text{HMLO}(t) + e_i(t)$$

The industry portfolios and factors are described in Tables 2.1 and 2.2. R_i is the portfolio's monthly excess return (in percent). R_M is the return on the CRSP value-weighted index minus the one-month T-bill rate. SMBO is the return on a portfolio of small stocks minus the return on a portfolio of big stocks, orthogonalized with respect to R_M . HMLO is the return on portfolio of high-B/M stocks minus the return on a portfolio of low-B/M stocks, again orthogonalized with respect to R_M . The table reports ordinary least squares estimates of the equations and the Gibbons, Ross, and Shanken (1989) F-test of the intercepts.

Portfolio	a		b		s		h		Adj R ²
	Coeff.	Std. err.	Coeff.	Std. err.	Coeff.	Std. err.	Coeff.	Std. err.	
Nat. resources	-0.05	0.20	0.97*	0.04	0.03	0.07	-0.02	0.07	0.57
Construction	-0.23*	0.10	1.14*	0.02	0.31*	0.04	0.17*	0.03	0.89
Food, tobacco	0.38*	0.12	0.90*	0.03	-0.12*	0.04	-0.01	0.04	0.77
Consumer products	-0.15	0.12	1.18*	0.03	0.68*	0.04	0.19*	0.04	0.86
Logging, paper	0.01	0.11	1.11*	0.02	0.05	0.04	0.05	0.04	0.85
Chemicals	0.21*	0.10	0.98*	0.02	-0.21*	0.04	-0.20*	0.03	0.85
Petroleum	0.29	0.19	0.81*	0.04	-0.45*	0.07	0.12*	0.07	0.53
Mach., equipment	0.04	0.10	1.11*	0.02	0.14*	0.04	-0.28*	0.04	0.87
Transportation	-0.24	0.12	1.08*	0.03	0.20*	0.04	0.28*	0.04	0.83
Utilities, telecom.	-0.06	0.10	0.65*	0.02	-0.26*	0.04	0.38*	0.04	0.74
Trade	0.03	0.14	1.13*	0.03	0.26*	0.05	0.01	0.05	0.80
Financial	-0.04	0.08	1.00*	0.02	-0.04	0.03	0.21*	0.03	0.89
Services, other	0.16	0.11	1.38*	0.03	0.74*	0.04	-0.17*	0.04	0.90
GRS F ^a	2.63								
(p-value)	(0.003)								

^a The GRS F-statistic equals $(T-N-K+1) / [N(T-K)] \cdot a' \Sigma^{-1} a$, where a is the vector of intercept estimates, Σ is the estimate of the covariance matrix of a , T is 368 (months), N is 13 (portfolios), and K is 4 (independent variables). Under the null hypothesis that all intercepts are zero, and assuming that returns are multivariate normal, this statistic is distributed as F (d.f. 13, 352).

* Denotes coefficients that are greater than two standard errors from zero.

be interpreted as the average factor sensitivities of the industries. Therefore, unless some industries were ‘distressed’ throughout the sample period, the significant explanatory power of SMBO and HMLO suggests that they proxy for more than just distress factors. Instead, the mimicking portfolios appear to reflect information relevant to a broad cross section of firms (see also Section 2.4.3).

The factors, however, cannot completely explain cross-sectional variation in average returns. Under the hypothesis that the three-factor model explains average returns, the intercepts in the time-series regressions should be zero. Table 2.5 shows that several intercepts are individually significant, and the Gibbons, Ross, and Shanken (1989) F-statistic rejects at the 1% level the restriction that all are zero. Economically, the intercepts are generally small, but two deviate from zero by over 3% annually. In sum, SMBO and HMLO proxy for pervasive risk factors in industry portfolios, and the three-factor model provides a reasonable, though not perfect, description of average returns.¹⁰

Table 2.6 reports SUR estimates of the conditional model. For simplicity, I do not report the constant terms of the intercepts and factor loadings (a_{i0} , b_{i0} , s_{i0} , and h_0). Since the industries’ B/M ratios are measured as deviations from their time-series means, the constant terms are simply estimates of the average coefficients, and they are nearly identical to the unconditional results in Table 2.5. Across all parameters, the mean absolute difference between the constant terms and the unconditional estimates in Table 2.5 is 0.017; for the intercepts only, it is 0.006. The similarity between the two sets of regressions indicates that changes in the loadings are largely uncorrelated with the factors.

The interactive terms with B/M are more interesting for our purposes. The table shows that B/M captures time-variation in risk, but does not appear to directly predict expected returns.

¹⁰ As a robustness check, I also estimate heteroskedastic-consistent standard errors and an asymptotically valid χ^2 statistic for the hypothesis that all intercepts are zero (based on the covariance estimates of White, 1984; see also Shanken, 1990). The results are not sensitive to heteroskedasticity adjustments.

Table 2.6

Conditional three-factor regressions: Industry portfolios

$$R_i = a_{i0} + a_{i1}B/M_i + (b_{i0} + b_{i1}B/M_i) R_M + (s_{i0} + s_{i1}B/M_i) SMBO + (h_{i0} + h_{i1}B/M_i) HMLO + e_i$$

The industries and factors are described in Tables 2.1 and 2.2. R_i is the portfolio's monthly excess return (%) and B/M_i is the natural log of its book-to-market ratio, as a deviation from its mean. R_M is the excess return on the CRSP value-weighted index. SMBO is the return on small stocks minus the return on big stocks, with $\text{cov}(R_M, SMBO) = 0$. HMLO is the return on high-B/M stocks minus the return on low-B/M stocks, with $\text{cov}(R_M, HMLO) = 0$. The table reports SUR estimates of the interactive terms, a_{i1} , b_{i1} , s_{i1} , and h_{i1} , which measure time-variation in the intercept and loadings.

Portfolio	a_1		b_1		s_1		h_1	
	Coeff.	Std. err.	Coeff.	Std. err.	Coeff.	Std. err.	Coeff.	Std. err.
Nat. resources	-0.13	0.62	-0.01	0.12	-0.28	0.23	-0.22	0.20
Construction	-0.36	0.24	-0.16*	0.05	0.05	0.09	0.14	0.09
Food, tobacco	-0.12	0.26	0.02	0.05	0.23*	0.09	0.40*	0.09
Consumer products	0.03	0.23	-0.17*	0.04	0.12	0.08	0.27*	0.07
Logging, paper	-0.28	0.38	-0.06	0.07	-0.11	0.14	0.34*	0.13
Chemicals	-0.30	0.28	-0.05	0.06	0.06	0.10	0.28*	0.10
Petroleum	0.62	0.54	-0.05	0.12	-0.32	0.20	-0.16	0.19
Mach., equipment	-0.72*	0.22	0.03	0.05	0.18*	0.08	0.29*	0.08
Transportation	0.07	0.30	-0.24*	0.07	0.00	0.12	0.02	0.10
Utilities, telecom.	0.33	0.26	0.01	0.06	0.12	0.10	0.04	0.09
Trade	-0.05	0.32	-0.02	0.07	0.27*	0.12	0.55*	0.10
Financial	0.43	0.31	0.07	0.07	-0.17	0.12	-0.33*	0.11
Services, other	-0.08	0.27	0.00	0.05	0.23*	0.10	0.35*	0.08
Average	-0.04		-0.05*		0.03		0.15*	
(std. err.)	(0.09)		(0.02)		(0.03)		(0.03)	
χ^2 ^a	14.27		38.94*		24.72*		95.76*	
(p-value)	(0.355)		(0.000)		(0.025)		(0.000)	

^a $\chi^2 = c' \Sigma^{-1} c$, where c is the vector of coefficient estimates and Σ is the estimate of the covariance matrix of c . Under the null that all coefficients are zero, this statistic is asymptotically distributed as χ^2 (d.f. 13).

* Denotes coefficients that are greater than two standard errors from zero or χ^2 statistics with a p-value less than 0.050.

The χ^2 statistics easily reject the hypotheses that B/M is unrelated to the loadings on R_M , SMBO, and HMLO. The B/M interactive terms with R_M , SMBO, and HMLO are over two standard errors from zero for 3 portfolios, 4 portfolios, and 8 portfolios, respectively. B/M tends to be positively related to the loadings on the size and book-to-market factors (financial firms are the exception), but negatively related to market betas. Interpreting increases in B/M as evidence of distress, it appears that market risk becomes relatively less important for distressed industries. While somewhat surprising, a similar result has been documented previously for firms near bankruptcy (e.g., McEnally and Todd, 1993).

In contrast, there is no evidence that B/M explains economically or statistically significant variation in the intercepts. None of the interactive terms with the intercepts is significantly positive, eight of the 13 estimates are negative, and the average coefficient, -0.04, is insignificantly different from zero (standard error of 0.09).¹¹ In fact, the only significant coefficient is actually negative (for the machinery and equipment industry), which is inconsistent with the overreaction story. In addition, the χ^2 statistic cannot reject the hypothesis that all coefficients on B/M are zero, with a p-value of 0.355. Thus, variation in risk appears to explain any association between B/M and expected returns.

Importantly, the lack of statistical significance is not driven by low power. The standard error of the average coefficient is relatively low, 0.09, and allows rejection of economically significant slopes. For example, suppose that the actual coefficient is two standard errors above the sample estimate, or 0.13. This coefficient maps into less than a 0.09% change in the monthly intercept when B/M varies by 0.66, twice its standard deviation for the typical portfolio. The OLS estimates (not reported) of the conditional regressions support these

¹¹ There is no mechanical reason that the average coefficient is zero. Conditional asset-pricing tests typically use the same conditioning variables for all portfolios, and some linear combination of the coefficients must be zero. However, no linear constraint is imposed on the coefficients here because B/M differs across portfolios. For example, aggregate B/M explains, on average, half of the variation in an industry's B/M ratio. In fact, when B/M is measured net of an aggregate index, the average correlation across portfolios is necessarily close to zero.

conclusions. Individually, the B/M coefficients are not significant, with an average estimate equal to -0.07 (standard error of 0.11), and the χ^2 statistic does not reject the joint restriction that all are zero (p-value of 0.370). The evidence is inconsistent with the argument that B/M proxies for mispricing in stock returns.

We saw earlier that the slope estimate is biased upward in a simple regression of returns on lagged B/M. The B/M term in the three-factor regression is likely to be biased upward as well, which would strengthen the conclusions above. An ad hoc estimate of the bias can be obtained by substituting the residuals from the three-factor regressions for the simple-regression error terms in eq. (2.8). The average bias estimated this way, 0.17, is much smaller than the bias in the simple regressions, 0.85. Bootstrap simulations like those described in Section 4 produce a similar estimate, 0.18.

2.4.2. Size and book-to-market portfolios

Tables 2.7 and 2.8 report similar findings for the size and book-to-market portfolios. In the unconditional regressions in Table 2.7, SMBO and HMLO capture significant co-movement in stock returns. For the size portfolios, the loadings on all factors are greatest for the smallest portfolios and decrease almost monotonically with size. They range from 0.91 to 1.22 on the market factor, -0.31 to 1.39 on SMBO, and -0.08 to 0.30 on HMLO. For the book-to-market portfolios, the loadings on SMBO and HMLO increase almost monotonically from the lowest to the highest deciles. The coefficients vary widely across portfolios. The cross-sectional spread is -0.09 to 0.87 for the loadings on SMBO and -0.77 to 0.97 for the loadings on HMLO. Market betas are generally close to one, ranging from 0.91 to 1.13, but are highest for the extreme portfolios (portfolios 1a and 10b). Consistent with the evidence in Fama and French (1993), the multivariate F-statistic rejects the asset-pricing restriction that all intercepts are zero. However, the deviations from zero are small (with the exception of low-B/M portfolio), and the three-factor model provides a fairly accurate description of average stock returns.

Table 2.7

Unconditional three-factor regressions: Size and book-to-market portfolios

$$R_i(t) = a_i + b_i R_M(t) + s_i \text{SMBO}(t) + h_i \text{HMLO}(t) + e_i(t)$$

The portfolios and factors are described in Tables 2.1 and 2.2. R_i is the portfolio's monthly excess return (in percent). R_M is the return on the CRSP value-weighted index minus the one-month T-bill rate. SMBO is the return on a portfolio of small stocks minus the return on a portfolio of big stocks, orthogonalized with respect to R_M . HMLO is the return on portfolio of high-B/M stocks minus the return on a portfolio of low-B/M stocks, again orthogonalized with respect to R_M . The table reports ordinary least squares estimates of the equations and the Gibbons, Ross, and Shanken (1989) F-test of the intercepts.

Portfolio	a		b		s		h		Adj R ²
	Coeff.	Std. err.	Coeff.	Std. err.	Coeff.	Std. err.	Coeff.	Std. err.	
<i>Panel A: Size portfolios</i>									
Smallest	-0.15	0.09	1.16*	0.02	1.39*	0.03	0.30*	0.03	0.93
2	-0.13*	0.05	1.19*	0.01	1.09*	0.02	0.19*	0.02	0.97
3	-0.11*	0.04	1.22*	0.01	0.96*	0.02	0.15*	0.01	0.98
4	-0.05	0.05	1.21*	0.01	0.84*	0.02	0.12*	0.02	0.98
5	0.03	0.05	1.19*	0.01	0.74*	0.02	0.11*	0.02	0.98
6	0.09	0.05	1.15*	0.01	0.59*	0.02	0.11*	0.02	0.97
7	-0.01	0.05	1.12*	0.01	0.43*	0.02	0.10*	0.02	0.97
8	0.02	0.05	1.12*	0.01	0.27*	0.02	0.13*	0.02	0.97
9a	0.01	0.06	1.07*	0.01	0.09*	0.02	0.14*	0.02	0.95
9b	-0.01	0.05	1.04*	0.01	0.05*	0.02	0.11*	0.02	0.96
10a	0.00	0.05	0.99*	0.01	-0.12*	0.02	0.03	0.02	0.96
Largest	0.04	0.04	0.91*	0.01	-0.31*	0.01	-0.08*	0.01	0.97
GRS F ^a	2.43								
(p-value)	(0.005)								

(Table 2.7 continued on next page)

Table 2.7. Continued.

Portfolio	a		b		s		h		Adj R ²
	Coeff.	Std. err.	Coeff.	Std. err.	Coeff.	Std. err.	Coeff.	Std. err.	
<i>Panel B: Book-to-market portfolios</i>									
Lowest	0.40*	0.09	1.12*	0.02	-0.02	0.03	-0.77*	0.03	0.91
1a	0.11	0.09	1.07*	0.02	-0.09*	0.03	-0.42*	0.03	0.90
2	-0.04	0.07	1.09*	0.01	-0.07*	0.02	-0.25*	0.02	0.94
3	-0.06	0.07	1.03*	0.02	-0.06*	0.03	-0.10*	0.02	0.92
4	-0.10	0.08	0.99*	0.02	-0.04	0.03	0.09*	0.03	0.90
5	-0.14	0.08	0.95*	0.02	-0.01	0.03	0.19*	0.03	0.89
6	-0.05	0.08	0.91*	0.02	-0.06*	0.03	0.39*	0.03	0.90
7	0.08	0.07	0.93*	0.02	-0.01*	0.03	0.48*	0.03	0.91
8	0.06	0.07	0.93*	0.02	0.10*	0.03	0.65*	0.03	0.91
9	0.14	0.08	1.01*	0.02	0.25*	0.03	0.71*	0.03	0.92
10a	0.02	0.12	1.12*	0.03	0.52*	0.04	0.81*	0.04	0.87
Highest	-0.12	0.15	1.13*	0.03	0.87*	0.06	0.97*	0.05	0.82
GRS F ^a	2.24								
(p-value)	(0.010)								

^a The GRS F-statistic equals $(T-N-K+1) / [N(T-K)] \cdot a' \Sigma^{-1} a$, where a is the vector of intercept estimates, Σ is the estimate of the covariance matrix of a , T is 368 (months), N is 12 (portfolios), and K is 4 (independent variables). Under the null hypothesis that all intercepts are zero, and assuming that returns are multivariate normal, this statistic is distributed as F (d.f. 12, 353).

* Denotes coefficients that are greater than two standard errors from zero.

The conditional three-factor regressions are more important for the current paper. Table 2.8 reports SUR estimates for the conditional model, in which intercepts and factor loadings vary linearly with lagged B/M. As before, the constant terms in the regressions are similar to the unconditional coefficients in Table 2.7, and I report only the interactive terms with B/M.

The evidence supports the conclusion that B/M captures significant variation in risk, but has little power to directly predict expected returns. For both sets of portfolios, the χ^2 statistics strongly reject, at the 0.001 level, the hypothesis that B/M is unrelated to the factor loadings. B/M displays a consistently positive relation to the loadings on the size and book-to-market factors. For the 24 portfolios shown in Table 2.8, 15 of the interactive terms with SMBO are greater than two standard errors above zero, and only one is significantly negative. Similarly, 16 of the coefficients on HMLO are significantly positive, and only one is significantly negative. The relation between B/M and markets betas is mixed. An increase in B/M predicts smaller betas for ten portfolios and larger betas for eight portfolios. Together with Table 2.6, the conditional regressions provide considerable evidence that B/M explains variation in risk.

Changes in risk absorb nearly all of B/M's predictive ability. The interactive terms with the intercepts are generally small and statistically insignificant. The average coefficient for the size portfolios is 0.00 (standard error of 0.05) and for the book-to-market portfolios is 0.03 (standard error of 0.07). Neither estimate is statistically different from zero, and we can reject economically significant coefficients. For example, true coefficients of 0.10 and 0.17 are two standard errors above the averages reported in Table 2.8. These coefficients map into 0.06% and 0.10% changes in monthly expected returns, respectively, when B/M varies by twice its standard error for the typical portfolio. The findings are striking given the significant explanatory power of B/M in simple regressions (see Table 2.4). By controlling for changes in risk, the average slopes on B/M decrease from 0.27 to 0.00 for the size portfolios and 1.02 to 0.03 for the book-to-market portfolios. B/M does not appear to have incremental explanatory power in predicting returns.

Table 2.8

Conditional three-factor regressions: Size and book-to-market portfolios

$$R_i = a_{i0} + a_{i1}B/M_i + (b_{i0} + b_{i1}B/M_i) R_M + (s_{i0} + s_{i1}B/M_i) SMBO + (h_{i0} + h_{i1}B/M_i) HMLO + e_i$$

The portfolios and factors are described in Tables 2.1 and 2.2. R_i is the portfolio's monthly excess return (in percent) and B/M_i is the natural log of the portfolio's book-to-market ratio, measured as a deviation from its time-series mean. R_M is the excess return on the CRSP value-weighted index. SMBO is the return on a portfolio of small stocks minus the return on a portfolio of big stocks, orthogonalized with respect to R_M . HMLO is the return on portfolio of high-B/M stocks minus the return on a portfolio of low-B/M stocks, again orthogonalized with respect to R_M . The table reports SUR estimates of the interactive terms, a_{i1} , b_{i1} , s_{i1} , and h_{i1} , which measure time-variation in the intercepts and factor loadings.

Portfolio	a_1		b_1		s_1		h_1	
	Coeff.	Std. err.	Coeff.	Std. err.	Coeff.	Std. err.	Coeff.	Std. err.
<i>Panel A: Size portfolios</i>								
Smallest	-0.10	0.24	-0.10*	0.04	-0.08	0.09	0.17*	0.07
2	0.09	0.15	-0.10*	0.03	0.05	0.05	0.25*	0.04
3	0.09	0.12	-0.06*	0.02	0.07	0.04	0.17*	0.04
4	0.06	0.12	-0.06*	0.02	0.11*	0.04	0.22*	0.04
5	-0.09	0.12	-0.09*	0.02	0.08*	0.04	0.21*	0.04
6	-0.10	0.13	-0.05*	0.02	0.19*	0.05	0.17*	0.04
7	-0.01	0.15	0.02	0.03	0.33*	0.05	0.16*	0.04
8	0.03	0.14	0.09*	0.03	0.19*	0.05	0.18*	0.04
9a	-0.16	0.17	0.10*	0.03	0.23*	0.06	0.18*	0.05
9b	0.07	0.15	0.08*	0.03	0.14*	0.05	0.07	0.05
10a	0.18	0.15	0.09*	0.03	0.07	0.06	-0.06	0.06
Largest	-0.07	0.08	-0.08*	0.02	-0.21*	0.03	-0.02	0.03
Average (std. err.)	0.00 (0.05)		-0.01 (0.01)		0.10* (0.02)		0.14* (0.01)	
χ^2 ^a (p-value)	7.67 (0.810)		78.67* (0.000)		81.49* (0.000)		126.28* (0.000)	

(Table 2.8 continued on next page)

Table 2.8. Continued.

Portfolio	a ₁		b ₁		s ₁		h ₁	
	Coeff.	Std. err.	Coeff.	Std. err.	Coeff.	Std. err.	Coeff.	Std. err.
<i>Panel B: Book-to-market portfolios</i>								
Lowest	-0.70*	0.25	0.21*	0.05	0.24*	0.09	-0.18*	0.09
1b	-0.79*	0.27	0.10	0.06	-0.01	0.10	0.17	0.09
2	-0.04	0.23	-0.03	0.05	0.21*	0.08	0.05	0.07
3	0.67*	0.24	-0.08	0.04	0.06	0.09	0.09	0.07
4	0.43	0.25	-0.11*	0.05	0.19*	0.09	0.19*	0.08
5	0.24	0.25	-0.13*	0.05	0.30*	0.09	0.25*	0.08
6	-0.07	0.25	-0.13*	0.05	0.05	0.09	0.43*	0.07
7	0.35	0.25	0.00	0.04	0.12	0.09	0.32*	0.07
8	0.17	0.25	-0.07	0.05	0.18*	0.09	0.33*	0.07
9	0.34	0.26	0.13*	0.05	0.35*	0.10	-0.02	0.08
10a	0.05	0.39	0.18*	0.07	0.33*	0.14	0.27*	0.11
Highest	-0.26	0.43	0.16*	0.08	0.68*	0.16	0.51*	0.12
Average (std. err.)	0.03 (0.07)		0.02 (0.01)		0.23* (0.03)		0.20* (0.02)	
χ^2 ^a (p-value)	24.32* (0.018)		42.08* (0.000)		103.26* (0.000)		146.74* (0.000)	

^a $\chi^2 = c' \Sigma^{-1} c$, where c is the vector of coefficient estimates and Σ is the estimate of the covariance matrix of c . Under the null that all coefficients are zero, this statistic is asymptotically distributed as χ^2 (d.f. 12).

* Denotes coefficients that are greater than two standard errors from zero or χ^2 statistics with a p-value less than 0.050.

Individually, the estimates for the size portfolios are small, and the χ^2 statistic cannot reject the hypothesis that all coefficients are zero, with a p-value of 0.810. The results for the book-to-market portfolios, however, provide some evidence of predictability: two coefficients are significantly negative (-0.70 and -0.79 for portfolios 1a and 1b) and one is significantly positive (0.67 for the portfolio 3). I discount the significance of the negative coefficients since they are inconsistent with both the efficient-market and overreaction stories. Also, the positive coefficient is the maximum estimate observed after searching over many coefficients, which provides an upward-biased estimate of the true maximum.¹² Overall, the picture that emerges from Tables 2.6 and 2.8 is that B/M contains substantial information about the riskiness of stock portfolios, but does not directly predict expected returns. There is virtually no support for the overreaction hypothesis.

2.4.3. *Industry-neutral HML*

Daniel and Titman (1997) argue that HML does not proxy for a separate risk factor in returns, but explains return covariation only because similar types of firms become mispriced at the same time. Their argument suggests that an industry's B/M ratio and its loading on HML will be related even under the mispricing story. By construction, HML invests in stocks with high B/M ratios. When an industry's B/M increases, HML becomes weighted toward firms in that industry and will, therefore, tend to covary more strongly with the industry return. In this case, time-varying factor loadings on HML might help explain mispricing related to B/M. To check whether the results for industry portfolios are driven by changes in the industry composition of HML, I replicate the three-factor regressions using an 'industry-neutral' book-to-market factor.

¹² Bonferroni confidence intervals provide a straightforward way to incorporate searching into statistical significance. Viewed in isolation, the estimate for decile 3 has a one-sided p-value of 0.002. Recognizing that the estimate is the maximum over 37 total portfolios, the Bonferroni upper bound on the p-value is 0.002×37 , or 0.083. See Johnson and Wichern (1982, p. 197).

As detailed in Appendix A, HML equals the return on a portfolio of high-B/M stocks minus the return on a portfolio of low-B/M stocks. I construct an industry-neutral factor, HML-N, in exactly the same way, except that stocks are sorted by their industry-adjusted B/M ratios, defined as the firm's B/M ratio minus the value-weighted average for all firms in their industry. The industries are defined for this purpose using the same classifications as the industry portfolios. By construction, then, the adjusted B/M ratios for firms in each industry are distributed around zero, so every industry should be represented approximately equally in the high- and low-B/M portfolios used to obtain HML-N.¹³

Empirically, the sorting procedure does not dramatically alter the book-to-market factor. HML-N has an average monthly return of 0.44% and a standard deviation of 2.42%, compared with 0.38% and 3.00%, respectively, for HML. The correlation between the two book-to-market factors, 0.87, is fairly high, which suggests that much of the variation in HML is unrelated to industry factors. In fact, part of the difference between HML and HML-N is caused by the difference in their market betas. The market beta of HML-N equals -0.03, significantly closer to zero than the -0.23 beta of HML. I also note that the sorting procedure affects SMB, since the size factor controls for differences in stocks' B/M ratios. The new size factor, which I continue to call SMB, has a mean return of 0.24% and a standard deviation of 2.64%, compared with 0.30% and 2.91% for Fama and French's (1993) size factor. The two size factors are almost perfectly correlated, with a sample correlation of 0.99. As before, I orthogonalize these factors with respect to the market return for the three-factor regressions.

Table 2.9 reports conditional regressions for the industry portfolios. For simplicity, the table reports only the coefficient estimates because the standard errors are close to those in Tables 2.5 and 2.6 (most differ by less than 0.01). The results are surprisingly similar to the findings for the Fama and French factors. Like HML, HML-N explains significant co-

¹³ As an alternative, I also divided the industry-adjusted B/M ratios by the standard deviation across firms in the industry. This modification does not affect the qualitative results.

Table 2.9

Three-factor regressions with industry-neutral HML: Industry portfolios

$$R_i = a_{i0} + a_{i1}B/M_i + (b_{i0} + b_{i1}B/M_i) R_M + (s_{i0} + s_{i1}B/M_i) SMBO + (h_{i0} + h_{i1}B/M_i) HML-N + e_i$$

The industries are described in Table 2.1. R_i is the portfolio's monthly excess return (%) and B/M_i is the natural log of its book-to-market ratio, as a deviation from its mean. R_M is the excess return on the CRSP value-weighted index. SMBO is a portfolio of small stocks minus big stocks. HML-N is a portfolio of high-B/M stocks minus low-B/M stocks. HML-N is constructed by sorting stocks based on their industry-adjusted B/M ratios (the firm's B/M ratio minus its industry average). The table reports SUR estimates, and the standard errors are similar to Tables 2.5 and 2.6.

Portfolio	Intercept		R_M		SMBO		HML-N	
	a_0	a_1	b_0	b_1	s_0	s_1	h_0	h_1
Nat. resources	-0.06	-0.19	0.98*	0.00	0.00	-0.28	0.05	-0.39
Construction	-0.27*	-0.42	1.15*	-0.15*	0.24*	-0.01	0.32*	0.15
Food, tobacco	0.38*	-0.06	0.90*	0.02	-0.10*	0.29*	-0.13*	0.29*
Consumer products	-0.16	-0.02	1.20*	-0.12*	0.67*	0.12	0.18*	0.26*
Logging, paper	-0.02	-0.32	1.12*	-0.04	0.01	-0.17	0.11*	0.41*
Chemicals	0.20*	-0.23	0.99*	-0.05	-0.19*	0.06	-0.22*	0.20
Petroleum	0.32	0.71	0.81*	-0.05	-0.50*	-0.36	0.10	-0.36
Mach., equipment	-0.01	-0.70*	1.13*	0.03	0.15*	0.13	-0.23*	0.38*
Transportation	-0.24*	-0.01	1.12*	-0.21*	0.14*	0.07	0.39*	-0.04
Utilities, telecom.	0.00	0.33	0.65*	-0.01	-0.27*	0.19	0.24*	0.08
Trade	0.01	-0.07	1.14*	0.02	0.26*	0.29*	-0.06	0.53*
Financial	-0.01	0.38	0.99*	0.07	-0.02	-0.12	0.18*	-0.34*
Services, other	0.16	-0.07	1.37*	0.04	0.78*	0.23*	-0.27*	0.41*
Average (std. err.)	0.02 (0.02)	-0.05 (0.08)	1.04* (0.01)	-0.03* (0.02)	0.09* (0.01)	0.03 (0.03)	0.05* (0.01)	0.12* (0.03)
χ^2 ^a (p-value)	36.42* (0.001)	13.52 (0.408)	59841.86* (0.000)	30.37* (0.004)	727.70* (0.000)	23.76* (0.033)	288.46* (0.000)	77.23* (0.000)

^a $\chi^2 = c' \Sigma^{-1} c$, where c is the vector of coefficient estimates and Σ is the estimate of the covariance matrix of c . Under the null that all coefficients are zero, this statistic is asymptotically distributed as χ^2 (d.f. 13).

* Denotes coefficients that are greater than two standard errors from zero or χ^2 statistics with a p-value less than 0.050.

movement in returns: ten of the 13 unconditional factor loadings are greater than two standard errors from zero, and the χ^2 statistic strongly rejects the hypothesis that all are zero. In addition, B/M captures significant time-variation in the factor loadings. Focusing on HML-N, seven of the 13 interactive terms are more than two standard errors from zero, and both the average coefficient (0.12, standard error of 0.03) and the χ^2 statistic (p-value less than 0.001) reject the hypothesis of constant risk. Again, B/M does not predict returns after controlling for changes in risk. None of the interactive terms with the intercept, a_1 , is significantly positive, and more than half of the estimates are negative. The average coefficient is also negative, and the χ^2 statistic cannot reject that all coefficients are zero.

These results say several interesting things about the book-to-market factor. First, HML (or HML-N) appears to capture a risk factor in returns that is unrelated to industry, contrary to the arguments of Daniel and Titman (1997). Neither the variation in HML, nor its covariation with industry returns, changes substantially when I control for changes in HML's industry composition. Second, HML appears to proxy for more than a distress factor in returns, unless some industries were distressed throughout the sample period. The cross-sectional spread of the unconditional factor loadings on HML is large (0.66 compared with 0.73 for market betas), and the variation across individual stocks is undoubtedly greater. Thus, HML contains information about a broad cross section of firms regardless of whether they are currently distressed. Finally, changes in the industry composition of HML do not drive changes in the industry portfolios' factor loadings. B/M continues to explain significant time-variation in risk after controlling for changes in HML's industry composition. Taken as a whole, the evidence supports the argument that B/M relates to a priced risk factor in returns.

2.5. Summary and conclusions

Previous studies find that B/M explains significant cross-sectional variation in average

returns. That finding implies that, at a fixed point in time, B/M conveys information about the firm's expected return relative to other stocks. This essay addresses a related question: For a given portfolio, does B/M contain information about the portfolio's expected return over time? The time-series analysis complements research on the predictability of stock returns at the aggregate level, and provides an alternative to cross-sectional tests of the risk- and characteristic-based asset-pricing stories.

The main empirical tests focus on industry portfolios. I find some evidence that an industry's B/M ratio predicts changes in its expected return, but the high variance of monthly returns reduces the precision of the estimates. The average, bias-adjusted coefficient on B/M, 0.58, is similar to the cross-sectional slope, 0.50, estimated by Fama and French (1992). The size and book-to-market portfolios produce more reliable evidence that B/M predicts returns. The results suggest that B/M tracks economically large changes in expected returns.

The conditional multifactor regressions indicate that B/M captures time-variation in risk, as measured by the Fama and French (1993) three-factor model. B/M tends to be positively related to the loadings on the size and book-to-market factors, but its relation to market betas is more difficult to characterize. The general impression conveyed by the conditional regressions is that market risk becomes relatively less important as a portfolio's B/M ratio increases. While it is beyond the scope of the current paper, understanding the economic reasons for the pattern of coefficients would provide additional insights into the connection between B/M and risk. I simply note here that the positive association between B/M and the loadings on HML does not seem to be driven by industry-related variation in the book-to-market factor.

After controlling for changes in risk, B/M contains little additional information about expected returns. Time-variation in the intercepts of the three-factor model measures the incremental explanatory power of B/M. For the industry portfolios, the average estimate has the opposite sign predicted by the overreaction story, and it is not significantly different from zero. Across the 13 portfolios, eight coefficients are negative and none are significantly positive at

conventional levels. Results for the size and book-to-market portfolios support these inferences: the average coefficients are indistinguishable from zero and roughly half the estimates are negative. The evidence for these portfolios is especially striking given B/M's strong predictive power when it is used alone in simple regressions. I have also replicated the tests in this paper using a firm's size in place of its B/M ratio, and find results qualitatively similar to those for B/M. In short, the three-factor model appears to explain time-varying expected returns better than a characteristic-based model.

To interpret the results, it is important to remember that we can always find some factor model to describe expected returns under both the efficient-market and mispricing stories (see, e.g., Roll, 1977; Shanken, 1987). The tests obtain economic meaning only when restrictions are imposed on the model. According to asset-pricing theory, the factors should capture pervasive risk in the economy related to investment opportunities or consumption. Under the mispricing view, it seems unlikely that the factors would explain, unconditionally, substantial covariation in returns. Many industries have large unconditional loadings on both the size and book-to-market factors, which provides some evidence that the factors proxy for priced risk in the economy.

Unfortunately, the case for rational pricing is not entirely satisfactory. This essay has been concerned primarily with changes in expected returns over time, not with their average levels. Consistent with the results of Fama and French (1993, 1997) and Daniel and Titman (1997), I find that the unconditional intercepts in the three-factor model are not zero. Thus, the model does not explain average returns. Just as important, the risk factors captured by the size and B/M mimicking portfolios have not been identified. The rational-pricing story will remain incomplete, and perhaps unconvincing, until we know more about the underlying risks.

On the Predictability of Stock Returns: Theory and Evidence

Chapter 3 Estimation risk, market efficiency, and the predictability of returns¹⁴

The analysis in Chapter 2 adopts the traditional perspective that predictability might arise either from mispricing or from changes in risk. As discussed in the introduction, market efficiency requires that prices ‘fully reflect all available information.’ To formalize this idea for empirical testing, Fama (1976) distinguishes between the probability distribution of returns perceived by ‘the market,’ based on whatever information investors view as relevant, and the true distribution of returns conditional on all information. The market is said to be informationally efficient if these distributions are the same. As an obvious consequence, market efficiency implies that investors correctly anticipate any cross-sectional or time-variation in true expected returns.

Market efficiency is closely related to the ‘rational expectations’ property analyzed by Muth (1961) and Lucas (1978). In Lucas’s model, asset prices are a function of the current level of output, whose behavior over time is known by investors. Consumers make investment decisions based, in part, on their expectations of future prices. Rational expectations requires that the pricing function implied by consumer behavior (the true pricing function) is the same as the pricing function on which decisions are based (the perceived pricing function). Lucas shows that rational expectations can, and generally will, give rise to predictable variation in expected returns (see also LeRoy, 1973). Intuitively, changes in economic conditions will lead to changes in the discount rate and, consequently, predictable returns. Thus, researchers must judge whether the empirical patterns in returns are consistent with credible models of rational behavior or can be better explained by irrational mispricing.

In this essay, we argue that there is a third potential source of return predictability:

¹⁴ This essay represents joint work with Jay Shanken.

estimation risk. In the asset-pricing literature, estimation risk refers to investor uncertainty about the true parameters of the return- or cashflow-generating process. Because investors do not know the true distribution, they must estimate the parameters using whatever information is available, which can be formally modeled using Bayesian analysis. The parameter uncertainty increases the perceived risk in the economy and necessarily influences portfolio decisions. As a result, estimation risk will affect equilibrium prices and expected returns. We show that, in equilibrium, estimation risk can be a source of predictability in a way that differs from other models with rational investors.

The theoretical literature typically focuses on the subjective distribution perceived by investors. The subjective distribution combines investors' prior beliefs with the information contained in observed data. This distribution represents investors' best guess about future returns or cashflows, and is therefore relevant for investment decisions.¹⁵ Our paper emphasizes instead the true distributions of prices and returns which arise endogenously in equilibrium. The true distribution simply refers to the actual, or observable, distribution from which prices or returns are drawn. Under the standard definition of market efficiency, the true and subjective distributions are the same. However, that definition goes well beyond the intuitive notion that prices fully reflect available information, and implicitly assumes that investors know the parameters of the cashflow process. In the presence of estimation risk, the two distributions necessarily differ since the true distribution depends on the unknown parameters. We should stress that 'true' does not mean 'exogenous': the true distribution of returns must be endogenous because prices clearly depend on investors' beliefs.

Our central result is easy to summarize: with estimation risk, the observable properties of prices and returns can differ significantly from the properties perceived by rational investors.

¹⁵ See Zellner (1971) and Berger (1985) for a general introduction to Bayesian analysis and Bawa, Brown, and Klein (1979) for an application to portfolio theory. Jobson, Korkie, and Ratti (1979), Jorion (1985), Kandel and Stambaugh (1996), Stambaugh (1998), and Barberis (1999) also discuss portfolio choice when investors must estimate expected returns.

For example, returns can appear predictable based on standard empirical tests even when they are not predictable by rational investors. The reason is simply that empirical tests estimate the true properties of returns, and these properties will typically differ from those under the subjective distribution. An example should help illustrate the point. Suppose dividends are normally distributed and independent over time with unknown mean δ and known variance σ^2 (in our parlance, this is the true distribution). From the investors' perspective, the mean of the dividend process is random, represented by a posterior belief about δ . However, for an empirical test, the process that generates actual dividends has a fixed, constant mean. The sampling distribution of any statistic calculated from dividends – say, an autocorrelation coefficient – depends only on this true distribution. In a similar way, the true distribution of *returns* is relevant for empirical tests even when it is unknown. To put the idea a bit differently, returns can be predictable under the true distribution, when they are not predictable by investors, since this distribution conditions on unknown information. We show that standard empirical tests, like predictive regressions and volatility tests, can in principle detect this predictability.

We develop these ideas in a simple overlapping-generations model of capital market equilibrium. Investors have imperfect knowledge about an exogenous dividend process, and they estimate the parameters based on current and past cashflows. For simplicity, we initially assume that all parameters are constant over time. We later extend the model to incorporate periodic shocks to the dividend process, in which case investors never fully learn the true distribution. Throughout, investors are assumed to be rational and use all available information when making decisions. As long as estimates of expected cashflows diverge from the true values, asset prices deviate from their values in the absence of estimation risk. However, prices tend to move toward these 'fundamental' values over time as investors update their beliefs. Through this process of updating, parameter uncertainty affects the predictability, volatility, and

cross-sectional distribution of returns.

The model shows that estimation risk can induce return behavior that resembles irrational mispricing. In our benchmark model without estimation risk, returns are unpredictable using past information. When investors must estimate the mean of the cashflow process, returns become predictable based on past dividends, prices, and returns. For example, when investors begin with a diffuse prior over the mean of the dividend process, stock prices appear to react too strongly to realized dividends, and returns become negatively related to past dividends and prices. In a fairly general sense, it appears that this phenomenon is inherent in a model with estimation risk because investors' 'mistakes' eventually reverse as they learn more about the underlying parameters. However, the predictability induced by estimation risk can take the form of either reversals or continuations (or neither), depending on investors' prior beliefs and on the underlying cashflow process (we discuss these issues further in Section 3.5). When investors have prior information about the dividend process, they may appear to react too slowly to new information, giving rise to momentum.

Predictability in the model is fundamentally different from predictability in other models with rational investors, such as that of Lucas (1978). The difference is illustrated most easily by considering the case of risk-neutral investors. In a model with perfect information, excess stock returns must be unpredictable if investors are risk-neutral. This does not have to be true with estimation risk. We show that excess stock returns can be predictable, under the true distribution, even with rational, risk-neutral investors. This predictability is consistent with rational expectations because investors do not know the true distribution. Nonetheless, the predictability can be detected by standard empirical tests. To reiterate our earlier point, excess returns remain unpredictable from the perspective of rational investors, but empirical tests estimate the true, not the subjective, distribution.

The example with risk-neutral investors shows that some basic properties of asset prices do not hold with estimation risk. Most importantly, investor rationality no longer implies that

return *surprises* must be uncorrelated with any element of investors' information set. In fact, return surprises will often be correlated with past prices if investors must estimate expected cashflows. The idea is simple. Suppose that prices equal the discounted present value of expected future dividends, assumed to be independent and identically distributed over time, and assume that investors do not know the mean of the dividend process. If a representative investor's estimate at a given point in time is, say, higher than the true mean, the price of the stock will be inflated above its 'fundamental' value. Furthermore, future dividends will be drawn from a true distribution with a lower mean than the market's estimate, and investors will, on average, perceive a negative surprise over the subsequent period. It follows that relatively high prices predict relatively low future returns.

This story resembles the standard mispricing argument, but with some important differences. Given estimation risk, the reversals are driven by completely rational behavior on the part of investors. The reversals arise precisely because prices *do* fully reflect all available information at each point in time. In fact, investors know that returns are negatively autocorrelated but cannot take advantage of it. They would want to exploit this pattern by investing more aggressively when the market's best estimate is less than the true mean of the dividend process, but of course they cannot know when this is the case. In contrast, DeLong, Shleifer, Summers, and Waldmann (1990), Daniel, Hirshleifer, and Subramanyam (1998), and Barberis, Shleifer, and Vishny (1998) generate return predictability by assuming irrationality on the part of investors. Investors misperceive the true return-generating process because of behavioral biases, not because they have imperfect information about returns.

The discussion has emphasized the time-series properties of returns. We also examine the cross section of expected returns. Curiously, for many years the conventional wisdom has been that estimation risk is largely irrelevant for equilibrium, although it is important for individual portfolio selection. For example, Bawa and Brown (1979) argue that estimation risk does not affect market betas or the expected return on the market portfolio. They conclude that

‘in empirical testing of equilibrium pricing models, one should not necessarily be concerned with the problem of estimation risk – or expect estimation risk to be a factor explaining any possible deviation between CAPM and observed market rates of returns,’ (p. 87).

More recently, Coles and Loewenstein (1988) argue that many of Bawa and Brown’s conclusions are driven by the questionable assumption that the return-generating process is exogenous. Coles and Loewenstein take end-of-period payoffs as exogenous, and allow prices and expected returns to adjust in equilibrium. They show that estimation risk affects fundamental economic features like relative prices, expected returns, and betas, although they continue to find that the CAPM holds in equilibrium.

Bawa and Brown (1979) and Coles and Loewenstein (1988) both examine the *subjective* distribution of returns. Its relevance for empirical research is questionable: although equilibrium imposes pricing restrictions under the subjective distribution, empirical tests use returns that are generated from the true distribution. Beliefs are relevant only insofar as they impact observable quantities. The basic distinction between the true and subjective distributions has typically been glossed over in the cross-sectional literature. Because the two distributions differ with estimation risk, we show that observed returns will typically deviate from the predictions of the CAPM, even when investors attempt to hold mean-variance efficient portfolios. Moreover, the deviations can be predictable, in either time-series or cross-sectional regressions, using past dividends and prices.

In short, our primary message is that estimation risk drives a wedge between the distribution perceived by investors and the distribution estimated by empirical tests. Although investors are rational, the empirical properties of prices and returns can look very different from the properties under the subjective distribution. Stock returns can appear predictable, in time-series or cross-sectionally, even though they are not from the perspective of rational investors. As a result, parameter uncertainty has important implications for characterizing and testing market efficiency. Our point here is not to argue that estimation risk necessarily explains

empirically-observed asset-pricing anomalies. Rather, we emphasize that many so-called ‘tests of market efficiency’ cannot distinguish between an efficient market with estimation risk and an irrational market. We believe that a world with estimation risk is the appropriate benchmark for evaluating apparent deviations from market efficiency.

Our results extend a growing literature on learning and parameter uncertainty. In the continuous-time literature, Merton (1971) and Williams (1977) show that parameter uncertainty creates a ‘new’ state variable representing investors’ current beliefs, and the hedging demand associated with this state variable can cause deviations from the CAPM (see also Detemple, 1986; Dothan and Feldman, 1986; Gennotte, 1986). Our results are different because investors in our model attempt to hold mean-variance efficient portfolios; it is their mistakes, not their hedging demands, that induce deviations from the CAPM. Stulz (1987) and Lewis (1989) also point out that prices can appear to overreact or underreact to information simply because investors must learn about the underlying true process. Wang (1993) and Brennan and Xia (1998) show that learning about an unobservable state variable might increase return volatility, but the effect on predictability is less clear. Finally, Timmermann (1993, 1996) recognizes that parameter uncertainty might induce both predictability and excess volatility. We extend his work by analyzing an equilibrium model with fully rational (Bayesian) investors, and we discuss market efficiency and the cross-section of expected returns.

The essay is organized as follows. Sections 3.1 and 3.2 introduce the basic model and derive capital market equilibrium. Section 3.3 examines the time-series properties of prices and returns and Section 3.4 explores the cross-sectional behavior of returns. Section 3.5 generalizes the model to incorporate informative priors, time-varying parameters, and non-stationary dividends, and presents simulation evidence from the general model. Section 3.6 concludes.

3.1. The model

We present a simple overlapping-generations model of capital market equilibrium in which

the dividend, or cashflow, process is taken as exogenous. Investors are uncertain about the true dividend process and update their beliefs with observed data. Many features of the model are borrowed from the economy analyzed by DeLong, Shleifer, Summers, and Waldmann (DSSW, 1990). Like DSSW, we examine capital market equilibrium when investors' beliefs diverge from the true distribution. In their model, noise traders' beliefs are exogenously specified and irrational. In contrast, investors in our model are rational and use all available information when making decisions.

3.1.1. Time

We analyze the properties of asset prices in an infinite-period model, $t = 1, \dots, \infty$. In single-period models of estimation risk, the end-of-period distribution of either returns or payoffs is exogenously specified (e.g., Bawa, Brown, and Klein, 1979; Coles and Loewenstein, 1988). In contrast, end-of-period prices in our model are determined by investors' beliefs, and both payoffs and returns are endogenous. When making decisions, investors must anticipate how market prices will react to the arrival of new information. Thus, the model permits a detailed investigation of both the time-series and cross-sectional behavior of returns.

3.1.2. Assets

We assume that there exists a riskless asset which pays real dividend r in every period. Following DSSW, the riskless asset is assumed to have perfectly elastic supply: it can be converted into, or created from, one unit of the consumption good in any period. As a result, its price in real terms must equal one and the riskless rate of return equals r .

The capital market also consists of N risky securities. As mentioned above, estimation risk has implications for both the time-series and cross-sectional behavior of asset prices. When we discuss the time-series properties of prices and returns, we examine a model with a single risky asset. The analysis with many risky assets focuses on the cross-sectional implications of

estimation risk.

Following Coles and Loewenstein (1988), we model investor uncertainty about an exogenously-specified cashflow process. Clearly, nothing can be learned about the return process if it is simply taken as exogenous, as assumed by Williams (1977) and Bawa and Brown (1979). If returns are endogenous, it is unclear how investors in the model would update their beliefs directly about the distribution of returns. For example, we doubt that any multiperiod model with estimation risk would produce returns that are independently and identically distributed (IID) over time. We show later that price reversals are inherent in a model with estimation risk, so it is unlikely that returns would be serially uncorrelated. Since the dividend process is assumed to be exogenous, we do not have to worry about how investors' beliefs affect its distribution.

The risky assets each have one unit outstanding and pay real dividend d_t , an $N \times 1$ vector, in period t . To develop the ideas in a simple framework, we initially assume that dividends are IID over time and have a multivariate normal distribution (MVN):

$$d_t \sim \text{MVN} [\delta, \Sigma], \quad (3.1)$$

where δ is the mean vector and Σ is a nonsingular covariance matrix. Notice that the parameters of this distribution are assumed to be constant over time. As a consequence, estimation risk will vanish as t goes to infinity. In reality, parameter uncertainty seems unlikely to disappear even after a long history of data. The economy evolves over time, and the underlying cashflow process undoubtedly changes as well. Therefore, we extend the model in Section 3.5 to include unobservable shocks to the true parameters which periodically renew estimation risk.

The IID assumption is not intended to be realistic, but dramatically simplifies the exposition. Again, we relax this assumption later and allow dividends to follow a geometric random walk. In addition, we have explored a model in which dividends are autocorrelated over time, and the qualitative results appear to be similar. Throughout the paper, investors are

assumed to know the form of the distribution function (IID and normal), but may not know its parameters.

3.1.3. *Investors*

Individuals live for two periods, with overlapping generations. Following DSSW, there is no first-period consumption, no labor supply decision, and no bequest. Therefore, in the first period individuals decide only how to invest their exogenously-given wealth. We assume that investors can be represented by a single agent with constant absolute risk aversion, or

$$U(w) = -\exp(-2\gamma w), \quad (3.2)$$

where w is second-period wealth and $\gamma > 0$ is the risk-aversion parameter.

Investors in this model do not have to allocate wealth across time. We ignore the intertemporal nature of the consumption problem and focus instead on estimation risk. It is almost immediate that investors will attempt to hold mean-variance efficient portfolios, and will not have hedging demands related to changes in investment opportunities (see Merton, 1973). This assumption limits the ways in which estimation risk can affect equilibrium, and distinguishes the predictability in our model from that in Merton (1971) and Williams (1977). In those papers, learning creates a state variable representing investors' beliefs, and the demand for risky assets contains a hedging component associated with this state variable. Our paper emphasizes a distinct phenomenon. We show that the difference between the true and subjective distributions can be a source of predictable returns.

The representative investor chooses a portfolio to maximize expected utility, where the expectation is taken over the investor's subjective belief about the distribution of next-period wealth. In all the cases we consider, both dividends and wealth are normally distributed. Consequently, it is easily shown that maximizing expected utility is equivalent to maximizing $\mu_w - \gamma\sigma_w^2$, where μ_w and σ_w^2 are the mean and variance of wealth. Let \mathbf{p} be the vector of

risky-asset prices and x_t be the vector of shares held in the portfolio. The investor will choose

$$x_t^* = \frac{1}{2\gamma} \left[\text{var}_t^s(p_{t+1} + d_{t+1}) \right]^{-1} \left[E_t^s(p_{t+1} + d_{t+1}) - (1+r)p_t \right], \quad (3.3)$$

where E_t^s and var_t^s denote the subjective expectation and variance at t .¹⁶ The first term in brackets is the covariance matrix of gross returns, and the second term is the expected excess gross return. Note that the optimal investment in the risky assets is not a function of initial wealth, an implication of constant absolute risk aversion. Also, given our assumptions that investors are short-lived and returns are multivariate normal, it is immediate that investors attempt to hold mean-variance efficient portfolios. Consequently, x_t^* is the Markowitz tangency portfolio under the subjective distribution.

Equilibrium in the economy, which treats current and future prices as endogenous, must satisfy eq. (3.3). In addition, equilibrium requires that the demands for the risky assets, given by x_t^* , equal their supply in every period. Setting $x_t^* = \mathbf{1}$, where $\mathbf{1}$ is an $N \times 1$ vector of ones, and solving for price yields

$$p_t = \frac{1}{1+r} \left[E_t^s(p_{t+1} + d_{t+1}) - 2\gamma \text{var}_t^s(p_{t+1} + d_{t+1}) \mathbf{1} \right]. \quad (3.4)$$

This equation gives the equilibrium current price in terms of next-period's price, which in turn will be endogenously determined.

3.2. Capital market equilibrium

This section derives capital market equilibrium with and without estimation risk. We assume throughout that investors correctly anticipate how prices will react to the arrival of new information. In other words, equilibrium satisfies the rational expectations property that the pricing function perceived by investors equals the true pricing function (Lucas, 1978). This

¹⁶ Throughout the paper we denote subjective moments with an 's' superscript.

condition does not imply, however, that investors' subjective belief about the *distribution* of returns equals the true *distribution*. Rational expectations, as we use the term, implies that these distributions are equal only if investors have perfect knowledge of the dividend process.

3.2.1. Equilibrium with perfect information

Suppose, initially, that investors know the dividend process. This equilibrium will serve as a convenient benchmark for the model with estimation risk. Since dividends are IID and the optimal investment in the risky asset does not depend on initial wealth, a natural equilibrium to look for is one in which prices are constant, or $p_t = p$. With constant prices, $E_t(p_{t+1} + d_{t+1}) = p + \delta$ and $\text{var}_t(p_{t+1} + d_{t+1}) = \Sigma$. Substituting into eq. (3.4) and solving for price yields

$$p = \frac{1}{r}\delta - \frac{2\gamma}{r}\Sigma\mathbf{1}. \quad (3.5)$$

The price of a risky asset equals its expected dividends discounted at the riskless rate minus a 'correction' for risk. Not surprisingly, an asset's contribution to the risk of the market portfolio (proportional to $\Sigma\mathbf{1}$; see below) is important, rather than its total variance. Investors require an expected rate of return that is higher than the riskless rate if the asset's 'market risk' is positive.

Many of the time-series implications of estimation risk can be investigated in a model with a single risky asset. The properties of this asset are identical to those of the market portfolio when there are many risky assets. In particular, the market portfolio M has weights proportional to the vector of prices, $p_t = p$. Its value, or price, equals

$$p_M = \mathbf{1}'p = \frac{1}{r}\delta_M - \frac{2\gamma}{r}\sigma_M^2, \quad (3.6)$$

where the dividend on the market portfolio has expectation $\delta_M = \mathbf{1}'\delta$ and variance $\sigma_M^2 = \mathbf{1}'\Sigma\mathbf{1}$. Since the variance is always positive, the expected return on the market portfolio is necessarily greater than the riskless rate. Referring back to the pricing function with many assets, it is straightforward to show that the general model collapses to eq. (3.6) when $N = 1$.

3.2.2. Equilibrium with estimation risk

The model above assumes that investors have perfect knowledge about the dividend process – that is, they know both the mean and the variance with certainty. We now relax this strong assumption. Specifically, suppose that investors begin with a diffuse prior over δ (the prior density function is proportional to a constant). Although this prior permits δ to be negative, it is the standard representation of ‘knowing little’ about the mean and simplifies the algebra. We later consider alternative prior beliefs. With an informative prior, investors assign less weight to the data and more weight to their initial beliefs, which can be important for the way prices behave in equilibrium. Consequently, the results in this and the next section should be interpreted as illustrative, but not completely representative, of the effects of estimation risk. For simplicity, we continue to assume that investors know the covariance matrix of dividends. Previous research finds that uncertainty about the covariance matrix is relatively unimportant (e.g., Coles, Loewenstein, and Suay, 1995), and we doubt that it would affect our basic conclusions.

Investors update their beliefs using Bayes rule, incorporating the information in observed dividends. With a diffuse prior, the posterior distribution of δ at time t is $MVN[\bar{d}_t, (1/t)\Sigma]$, where \bar{d}_t is the vector of average dividends observed up to time t . The subjective, or in Bayesian terms ‘predictive,’ distribution of dividends is

$$d_{t+1} \sim^s MVN\left[\bar{d}_t, \frac{t+1}{t}\Sigma\right]. \quad (3.7)$$

An investor’s best guess about the mean of the dividend process is simply the average realized dividend. The covariance matrix of the predictive distribution reflects both the true variance, Σ , and uncertainty about the mean, Σ / t .

From eq. (3.7), it is clear that the subjective distribution of dividends, and consequently future prices, differs from the true distribution. Rational expectations requires, however, that

investors correctly anticipate how equilibrium prices will be determined next period. We impose this requirement by recursively substituting for p_{t+k} in eq. (3.4), yielding¹⁷

$$p_t = \frac{1}{r} \bar{d}_t - 2\gamma \left[\sum_{k=1}^{\infty} \frac{1}{(1+r)^k} E_t^s \text{var}_{t+k-1}^s (p_{t+k} + d_{t+k}) \right] \cdot \iota. \quad (3.8)$$

Price is a function of expected dividends and the expected conditional variance of gross returns. Since estimation risk ‘scales up’ the predictive variance by $(t+1)/t$, the conditional variance of returns is unlikely to be constant. However, if price is a linear function of \bar{d}_t , then the conditional variance of returns will be a deterministic function of time. We look for an equilibrium that has this property.

If the conditional variance of returns is deterministic, then we can drop the expectations operator from the infinite sum in eq. (3.8). Variation in prices is driven entirely by the first term. Therefore, the subjective variance of returns is

$$\text{var}_t^s (p_{t+1} + d_{t+1}) = \left[1 + \frac{1}{r(t+1)} \right]^2 \left(\frac{t+1}{t} \right) \Sigma. \quad (3.9)$$

Substituting into eq. (3.8) yields the equilibrium pricing function:

$$p_t = \frac{1}{r} \bar{d}_t - 2\gamma f(t) \Sigma \iota, \quad (3.10)$$

where

$$f(t) = \sum_{k=1}^{\infty} \frac{1}{(1+r)^k} \left[1 + \frac{1}{r(t+k)} \right]^2 \left(\frac{t+k}{t+k-1} \right). \quad (3.11)$$

The equilibrium price is similar to the price with perfect information (eq. 3.5). The mean of the predictive distribution, \bar{d}_t , replaces the true mean in the first term and the function $f(t)$ replaces $1/r$ in the second term. The function $f(t)$ decreases as t gets larger and converges to $1/r$ in the

¹⁷ Eq. (3.8) imposes the transversality condition $\lim_{k \rightarrow \infty} E_t[p_{t+k}]/(1+r)^k = 0$, which will be satisfied in equilibrium.

limit. Since the probability limit of \bar{d}_t is δ , the equilibrium price with estimation risk converges to the price with perfect information. This is intuitive because, as mentioned above, estimation risk vanishes in the limit. In Section 3.5 we allow the true parameters to change, so that investors never completely learn the dividend process.

We noted in Section 3.1 that investors attempt to hold the tangency portfolio, which implies that the CAPM must describe expected returns under the subjective distribution. We will discuss the CAPM in more detail below, but for now we note that the market portfolio's value, or price, is

$$p_{M,t} = \frac{1}{r} p_t = \frac{1}{r} \bar{d}_{M,t} - 2\gamma f(t) \sigma_M^2, \quad (3.12)$$

where $\bar{d}_{M,t} = \frac{1}{t} \sum_{s=1}^t \bar{d}_s$ is the average dividend on the market portfolio from $t = 1$ to t . Referring back to the pricing function with many assets, it is straightforward to show that the general model collapses to eq. (3.12) when $N = 1$.

Several colleagues have noted that the pricing function in eq. (3.10) could also be generated by a model with a nonstationary dividend process and no estimation risk. In particular, suppose investors have perfect information and the *true* mean of the dividend process evolves over time as a function of average realized dividends (that is, $\delta_{t+1} = \bar{d}_t$). In this case, the pricing function would be identical to the price in our model. Notice, however, that our model should be distinguishable from one with nonstationary dividends. Prices and expected returns evolve quite differently in the two models. With a changing dividend process and perfect information, expected gross returns would be positively related to lagged dividends, and prices would exhibit no tendency to revert to a long-run mean. The opposite is true in our model; true expected returns are negatively related to lagged dividends and price fluctuations are temporary. Further, nonstationary dividends would not generate deviations from the CAPM.

3.3. The time-series properties of prices and returns

Equilibrium, derived above, is determined by the subjective distribution of returns. However, empirical tests use prices and returns drawn from the true distribution. As we emphasized before, the subjective and true distributions differ when there is estimation risk, even though investors know the true pricing function. In this section, we examine the time-series properties of prices and returns, highlighting the impact of estimation risk on market efficiency. The analysis considers a model with a single risky asset, interpreted as the market portfolio. In this case, the price of the risky asset is given by eqs. (3.6) and (3.12). We drop the subscript 'M' throughout this section for convenience.

In the model with perfect information, prices are constant and returns simply equal realized dividends. With estimation risk, prices fluctuate as investors update their beliefs about the dividend process. From the previous section, the change in price from t to $t+1$ equals

$$p_{t+1} - p_t = \frac{1}{r}(\bar{d}_{t+1} - \bar{d}_t) + 2\gamma[f(t) - f(t+1)]. \quad (3.13)$$

The change in price contains two components. The first term is random and reflects changes in investors' beliefs about expected dividends. The second term is deterministic and arises because estimation risk declines steadily over time. Since $f(t+1) < f(t)$, this component tends to make prices increase over time. When we talk about predictability, the deterministic portion serves only to add an additional, non-random component to the equations. Therefore, to focus on the main ideas, we assume in this section that investors are risk-neutral ($\gamma = 0$), causing the second term in the equation to drop out. None of the results are sensitive to this assumption.

3.3.1. Predictability

Previous studies argue that returns might be predictable either because business conditions change over time or because investors are irrational. However, these stories cannot explain why returns would be predictable in our model. The riskless rate, preferences, and the distribution of

cashflows do not change, so ‘business conditions’ are constant by construction. In addition, investors are rational and use all available information when making decisions, so irrational mispricing does not exist. In our model, estimation risk is the only source of predictability.

As noted above, returns equal dividends when investors have perfect information. With estimation risk, returns at $t+1$ equal

$$R_{t+1} = d_{t+1} + \frac{1}{r(t+1)} (d_{t+1} - \bar{d}_t). \quad (3.14)$$

The first term equals realized dividends, and the second term equals the change in price. At time t , investors’ best guess about dividends is given by \bar{d}_t ; when realized dividends differ from this expectation, investors revise their beliefs about the mean of the dividend process, which in turn affects prices.

Under the subjective distribution, it is clear that prices follow a martingale:

$$E_t^s[p_{t+1} - p_t] = 0. \quad (3.15)$$

However, the empirical properties of returns will differ from the perceived properties. The reason is simple. From the investor’s perspective, the expected dividend is random, represented by a posterior belief over δ . In contrast, for an empirical test, the dividend mean is fixed and constant, equal to whatever the true value actually is; the process that generates observed dividends does not have a random mean. Put differently, the observable properties of returns are conditional on the true dividend process even though it is unknown. Because of this fundamental difference between the true and subjective distributions, changes in prices can appear predictable to a researcher. From eq. (3.14), the true conditional expected return is

$$E_t[R_{t+1}] = \delta + \frac{1}{r(t+1)} (\delta - \bar{d}_t). \quad (3.15)$$

It is clear that R_{t+1} is negatively related to past dividends.¹⁸ Although dividends are IID by

¹⁸ For simplicity, we examine the predictability of *gross* returns rather than *rates of* return. The

assumption, price revisions are negatively correlated with past cashflows. The intuition is fairly straightforward. Prices depend on investors' best guess about future dividends, given by \bar{d}_t . The higher that past dividends have been, the lower that changes in beliefs are expected to be. As a result, price revisions move opposite to past cashflows.

From eq. (3.15), prices, dividends, and returns all predict time-variation in expected returns. For example, suppose we are interested in the autocovariance of returns:¹⁹

$$\text{cov}[R_t, R_{t+1}] = -\frac{1}{r t (t+1)} \sigma^2. \quad (3.16)$$

With estimation risk, returns are negatively autocorrelated. A researcher who ignores estimation risk, and observes that business conditions do not change, would come to the incorrect conclusion that investors overreact: higher returns today predict lower future returns. Similarly,

$$\text{cov}[d_t, R_{t+1}] = -\frac{1}{r t (t+1)} \sigma^2. \quad (3.17)$$

A high dividend today predicts lower future returns, which would suggest that investors naively extrapolate recent dividend performance into the future. However, investors are completely rational in our model and the predictability is driven entirely by estimation risk. Investors appropriately incorporate all relevant information, but today's dividend causes a revision in prices that moves opposite to expected returns.

We later present simulation evidence to show how estimation risk can affect empirical tests. To illustrate the results, Fig. 3.1 depicts a sample price path for the risky asset. The figure assumes that investors are risk-neutral and the riskless rate of return is 0.05. Dividends have mean 0.05 and standard deviation 0.10, taken to be similar to the dividend yield and

analysis with rates of return is more difficult because it involves expectations of ratios, but the qualitative results are similar.

¹⁹ This covariance is time-dependent because estimation risk declines over time. We will break the strong connection between time and predictability in Section 3.5 when we allow the true dividend process to change.

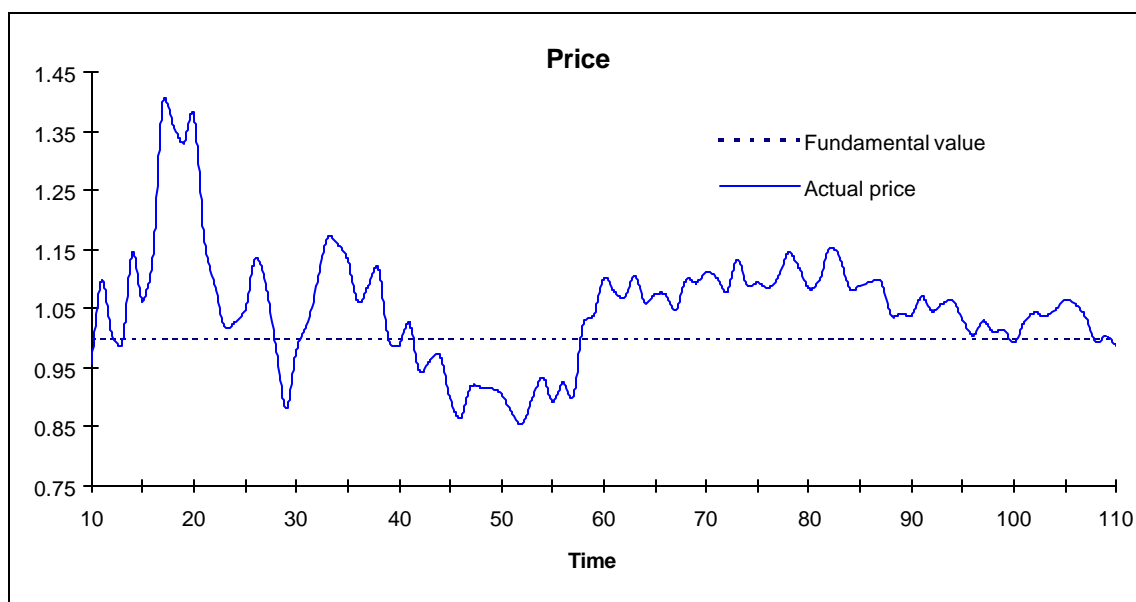


Figure 3.1
Equilibrium price of the risky asset

This figure illustrates a typical price path for the risky asset when the dividend process is known ('fundamental value' in the figure; see eq. 3.6 in the text) and when investors must estimate expected dividends ('actual price'; see eq. 3.12 in the text). The riskless rate is 0.05, dividends have true mean 0.05 and standard deviation 0.10, and investors are risk-neutral. Without estimation risk, the price of the risky asset is one. With estimation risk, the price depends on average dividends, which we randomly select from a Normal distribution.

volatility of dividends on the market portfolio. Under these assumptions, the price of the risky asset without estimation risk equals one and its expected rate of return is 0.05. The price with estimation risk depends on realized average dividends, which we randomly draw from a normal distribution. The figure shows that the price of the risky asset tends to revert towards 'fundamental' value. The sample autocorrelation in returns equals -0.10 for the periods shown ($t = 10$ to 110) and the correlation between rates of return and lagged prices equals -0.28. True conditional expected rates of return vary from 2.0% to 6.2%.²⁰

In this example, the mean-reversion in asset prices is obvious from the figure. The price-

²⁰ The example is for illustration purposes only. The reported statistics do not adjust for small-sample bias in the correlation and regression coefficients. We present more extensive simulation evidence in Section 3.5.

reversal effect of estimation risk might be observable to a researcher, yet prices at every point in time are set rationally. Investors ignore the negative relation between returns and dividends because it provides no useful information about *future* expected returns. In addition, the example suggests that the effects of parameter uncertainty can be relatively large, as measured by the variation in expected returns. Similar results for actual stock market data would be interpreted as relatively strong evidence against efficient markets. However, *ex ante*, investors in this example could not have forecast *any* variation in expected returns.

The analysis above considers the predictability of one-period returns. Investor expectations are highly persistent, however, and price reversals can take many periods to occur. As a result, the negative relation between returns and past dividends becomes stronger for long-horizon returns. Define the H-period return ending at $t+H$ as the sum of one-period returns, or $R_{t+H}^H = R_{t+1} + \dots + R_{t+H}$. Then the conditional expected H-period return is

$$E_t[R_{t+H}^H] = H\delta + \frac{H}{r(t+H)}(\delta - \bar{d}_t). \quad (3.18)$$

Similar to one-period returns, R_{t+H}^H is negatively related to past prices. Except for the substitution of $t+H$ for $t+1$ in the denominator, the expected return is H-times more sensitive to changes in average dividends than one-period returns. As a result, the price-reversal effect of estimation risk will be more pronounced in long-horizon returns. For example, the autocovariance of H-period returns is

$$\text{cov}[R_t^H, R_{t+H}^H] = -\frac{H^2}{r t(t+H)}\sigma^2, \quad (3.19)$$

which increases by a factor of H^2 as the horizon is lengthened (except for the change from $t+1$ to $t+H$ in the denominator). The variance of returns increases at a rate less than H , so returns become more negatively autocorrelated as the return horizon lengthens. Results are similar for the relation between expected returns and lagged dividends.

In short, estimation risk can be a source of predictability. However, the predictability of total returns does not say anything directly about market efficiency. In the model analyzed by Lucas (1978), for example, returns are predictable yet the market is efficient. To get a clearer picture of market efficiency, we need to examine the predictability of return *surprises*. A standard result in finance is that forecast errors should be unpredictable if investors are rational. Indeed, tests of market efficiency, like those analyzed by Shiller (1981) and Abel and Mishkin (1983), rely on the assumption that rational forecast errors are uncorrelated with past information. In the presence of estimation risk, we show that rationality no longer imposes this restriction. Investors form expectations based on past information, so forecast errors will be correlated under the true distribution with past cashflows.

The unexpected portion of returns, UR_{t+1} , is given by the difference between realized returns and investors' subjective expectation, or $R_{t+1} - E_t^s[R_{t+1}]$. In this section, we have assumed that investors are risk-neutral, implying that unexpected returns equal excess returns. From eq. (3.14),

$$UR_{t+1} = \left[1 + \frac{1}{r(t+1)} \right] (d_{t+1} - \bar{d}_t). \quad (3.20)$$

It follows that

$$E_t[UR_{t+1}] = \left[1 + \frac{1}{r(t+1)} \right] (\delta - \bar{d}_t). \quad (3.21)$$

Therefore, like total returns, the unexpected portion of returns is predictable based on past dividends, returns, and prices. It is precisely this result that differentiates predictability in our model from predictability in other models with rational investors. With perfect information, excess returns must be unpredictable if investors are risk-neutral. In contrast, once we allow for parameter uncertainty, excess returns can be predictable even with rational, risk-neutral investors.

Thus, not only do subjective expectations differ from true expectations, but they do so in a way that is predictable with prices and dividends. With incomplete information, investors form expectations based on observed dividends. If these have been, say, abnormally high, then price will be inflated above its fundamental (perfect information) value. Consequently, prices are related to future returns in a way that resembles overreaction. The predictability is consistent with rational expectations because it is based on the unknown, true distribution. We emphasize, however, that the true distribution determines the empirical properties of returns even though it is unknown. At the time portfolio decisions are made, investors cannot know whether past dividends have been above or below the true mean. Over time, investors learn more about expected cash flows and, looking back, can observe the negative relation between prices and unexpected returns (as illustrated by Fig. 3.1).

Throughout this section, we have found that parameter uncertainty creates price reversals and negative autocorrelation in returns. These results are relevant for the large empirical literature on excess volatility and apparent overreaction. However, several studies also document momentum in stock returns. Jegadeesh and Titman (1993), for example, find that short-term ‘winners’ (stocks that performed well over the past 3 to 12 months) have higher future returns than short-term ‘losers.’ In Section 3.5, we show that informative priors might give rise to momentum in returns. In addition, alternative cashflow processes, such as autocorrelated dividends, could generate momentum if investors are uncertain about the persistence of cashflows.

3.3.2. Price volatility

Price volatility is closely related to predictability (see, e.g., Campbell, 1991). For example, investor overreaction generally implies that returns will be both negatively autocorrelated and excessively volatile. Given our results above, it is clear that estimation risk will significantly affect the variance of prices and returns.

In the model without estimation risk, the variance of returns simply equals the variance of dividends, σ^2 . With parameter uncertainty, prices fluctuate over time as investors update their beliefs about the dividend process. In particular, the (true) conditional variance of price is

$$\text{var}_t[p_{t+1}] = \left[\frac{1}{r(t+1)} \right]^2 \sigma^2, \quad (3.22)$$

and the unconditional variance is

$$\text{var}[p_{t+1}] = \frac{1}{r^2(t+1)} \sigma^2. \quad (3.23)$$

Estimation risk increases both the conditional and unconditional variances of observed prices. Similar to inferences about return predictability, ignoring the effects of estimation risk would suggest investor overreaction. However, ‘excess’ volatility simply reflects parameter uncertainty; volatility is high precisely because investors rationally update their beliefs.

In the model, a relatively small amount of parameter uncertainty will substantially increase price volatility. Suppose, for example, that investors are risk-neutral, the riskless rate of return of 0.05, and dividends are distributed with mean 0.05 and standard deviation 0.10. (These are the values used in Fig. 3.1.) In this case, the value of the risky asset equals one when the dividend process is known. With parameter uncertainty, the standard deviation of p equals $2/\sqrt{t}$, which remains significant as a percentage of fundamental value for rather large t . When t is, say, 100 the standard deviation of price is 0.20. This implies that the length of a two-standard-deviation confidence interval is 80% of fundamental value, despite the fact that the subjective standard deviation of dividends is less than one percent greater than the true standard deviation.

Thus, the model suggests that prices might vary considerably around their ‘true’ values. The deviations are eventually reversed, giving rise to predictable variation in returns, yet investors are completely rational. Put differently, stock price movements do not have to be

explained by subsequent changes in dividends. Indeed, in our model, prices are *completely* uncorrelated with future dividends. Prices are backward looking and, ignoring estimation risk, investors appear to overreact to past information.

Asset prices can also violate the volatility bounds that have been the focus of much empirical research. For example, Shiller (1981) argues that an immediate consequence of ‘optimal forecasts’ is that

$$\text{var}(p_t) \leq \text{var}(p_t^*), \quad (3.24)$$

where p_t^* is the ex post rational price, or the price based on realized, rather than expected, dividends. That is, p_t^* is given by

$$p_t^* = \sum_{k=1}^{\infty} \frac{1}{(1+r)^k} d_{t+k} - 2\gamma f(t) \sigma^2. \quad (3.25)$$

With perfect information and rational investors, the bound holds because p_t^* equals the actual price plus a random, unpredictable forecast error. We saw above, however, that the forecast error with parameter uncertainty can be negatively related to price. In the current model, the variance of p_t^* is

$$\text{var}(p_t^*) = \sum_{k=1}^{\infty} \frac{1}{(1+r)^{2k}} \sigma^2 = \frac{1}{r^2 + 2r} \sigma^2. \quad (3.26)$$

Comparing this to eq. (3.23), we see that the volatility bound will be violated for $t \leq 1 + 2/r$. Perhaps more directly, however, prices violate the basic premise of the volatility-bound literature, that revisions in prices should only reflect changes in true expected dividends. With estimation risk, new information about future dividends does not have to correspond to changes in the true distribution. Thus, the volatility literature tests the joint hypothesis that investors are rational and have perfect information about the dividend process. Assuming that investors have less than perfect knowledge, it might be more surprising if prices did not violate the bounds.

The volatility of returns provides additional insights into the effects of estimation risk. The conditional variance of returns is

$$\text{var}_t[R_{t+1}] = \left[1 + \frac{1}{r(t+1)}\right]^2 \sigma^2. \quad (3.27)$$

Comparing this to the variance of the subjective distribution (eq. 3.9), we find the standard result that the subjective variance equals the true variance multiplied by $(t+1)/t$. Also, compared to the variance of returns when δ is known, which is just σ^2 , we see that estimation risk greatly increases return volatility if t is small. For example, if the riskfree rate is 0.05 and t is 50, the conditional variance of returns is roughly twice as big with estimation risk than without.

3.4. The cross-section of expected returns

We now return to the model with many assets and analyze the cross-sectional behavior of returns. In single-period models, Bawa and Brown (1979) and Coles and Loewenstein (1988) find that the CAPM continues to hold with estimation risk. These studies focus exclusively on the subjective distribution of returns, and find that estimation risk is largely irrelevant for equilibrium. We emphasize instead the observable behavior of prices and returns.

Before continuing, we should mention again that investors are assumed to begin with a diffuse prior. This assumption will affect the results in a variety of ways. For example, an informative prior can contain more information about some securities than others. The diffuse prior, on the other hand, is ‘symmetric.’²⁰ To see why this important, recall that unexpected returns in the model equal

$$\text{UR}_{t+1} = \left[1 + \frac{1}{r(t+1)}\right] (d_{t+1} - \bar{d}_t), \quad (3.28)$$

²⁰ With a symmetric prior, the prior covariance matrix is proportional to the true covariance matrix. That is, the prior distribution over δ has the form $\text{MVN}[\delta^*, \Sigma/h]$, where h is a measure of prior information. The diffuse prior can be interpreted as the limiting distribution as h approaches zero.

where UR_{t+1} , d_{t+1} , and \bar{d}_t are now $N \times 1$ vectors. In the brackets, the term ‘ $1/r(t+1)$ ’ gives the effect that unexpected dividends have on prices. With a diffuse prior, price revisions are proportional to the vector of unexpected dividends. This result will not generally hold with an informative prior. Dividends on assets with relatively high amounts of prior information provide clues about the values of other securities (Clarkson, Guedes, and Thompson, 1996; Stambaugh, 1997). We discuss informative priors further in Section 3.5.

3.4.1. Covariances and betas

When we talk about the CAPM, it will be useful to have a few results on covariances and market betas. With one risky asset, we saw that parameter uncertainty increases both the subjective and true volatility of returns. Similarly, with many assets, estimation risk scales up the true covariance matrix. In particular, the conditional covariance matrix of gross returns is

$$\text{var}_t [R_{t+1}] = \left[1 + \frac{1}{r(t+1)} \right]^2 \Sigma. \quad (3.29)$$

The effect of estimation risk is analogous to the single asset case: uncertainty about δ increases the true volatility of prices and returns. In the model with perfect information, prices are constant and the covariance matrix of returns simply equals Σ . With estimation risk, investor uncertainty increases all variances and covariances proportionally. Further, this statement describes both the subjective and true covariance matrices. Comparing eqs. (3.9) and (3.29), we find that the subjective covariance matrix equals the true covariance matrix multiplied by $(t+1)/t$, and both are proportional to Σ .²¹

Because estimation risk simply scales up the covariance matrix, it does not affect market betas (for gross returns). The market return is the sum of the asset returns, and market volatility

²¹ The assumption that investors begin with a diffuse prior is important here. With informative priors, the subjective and true covariance matrices may not be proportional, and neither has to be proportional to Σ . See Section 3.5.

increases by the same factor as the covariance matrix. Consequently, with and without parameter uncertainty, betas equal

$$\beta = \frac{1}{\text{var}(R_{M,t})} \text{cov}(R_t, R_{M,t}) = \frac{1}{\mathbf{1}' \Sigma_t} \Sigma_t \mathbf{1}. \quad (3.30)$$

Note also that eq. (3.30) gives both subjective and true market betas, which are the same because the two covariance matrices are proportional. Again, this result is an artifact of the diffuse prior. With an informative prior, subjective and true betas will not necessarily be the same, nor will they equal the betas without estimation risk.²²

3.4.2. Expected returns and the CAPM

In Section 3.3, we found that total and unexpected returns are predictable with lagged dividends and prices. With many assets, we consider instead deviations from the CAPM. We examine both the time-series and cross-sectional predictability of these deviations.

With and without estimation risk, the subjective distribution of returns is multivariate normal. Together with the assumption that investors derive utility only from end-of-period wealth, this implies that the CAPM must hold *under the subjective distribution*. In terms of gross returns, the CAPM says that

$$E_t^s [R_{t+1}] = r p_t + \beta [E_t^s (R_{M,t+1}) - r p_{M,t}]. \quad (3.31)$$

Eq. (3.31) can be verified by substituting for equilibrium price and subjective expected returns, derived above. Investors attempt to hold mean-variance efficient portfolios, which imposes the CAPM restriction on subjective expected returns. However, empirical tests use returns taken from the true, not the subjective, distribution.

To analyze the cross-sectional properties of returns, we focus on ex post deviations from

²² The analysis here focuses on *gross* returns, not *rates of* return. As noted by Coles and Loewenstein (1988), estimation risk does affect rate-of-return betas. Asset i 's rate-of-return beta equals its gross-return beta multiplied by p_{nt} / p_{it} . Rate-of-return betas will change unless relative prices remain the same, which will not be true in general.

the CAPM, given by

$$a_{t+1} = R_{t+1} - r p_t - \beta [R_{M,t+1} - r p_{M,t}]. \quad (3.32)$$

Note that a_{t+1} is similar to the vector of unexpected returns, except that the realized return on the market enters eq. (3.32) rather than the expected market return. We know from Section 3.3 that the market return is predictable based on past information. By examining a_{t+1} , rather than unexpected returns, we eliminate predictability that is related to the aggregate market.

Deviations from the CAPM must be unpredictable under the subjective distribution:

$$E_t^s[a_{t+1}] = 0. \quad (3.33)$$

In the absence of estimation risk, market efficiency implies that the true conditional expectation of a_{t+1} is zero. This restriction, of course, forms the basis for empirical tests of the CAPM. For example, cross-sectional regressions, like those in Fama and MacBeth (1973), indirectly test whether firm characteristics predict cross-sectional variation in $a_{i,t+1}$. The multivariate F-statistic of Gibbons, Ross, and Shanken (1989) tests whether the unconditional expectation of a_{t+1} is zero, which follows from the law of iterated expectations. Finally, various conditional asset-pricing tests directly examine the conditional expectation of a_{t+1} (e.g., Harvey, 1989; Shanken, 1990).

With parameter uncertainty, rational expectations no longer requires that the true expectation of a_{t+1} equals zero. Substituting for prices and returns in eq. (3.32) and taking expectations yields

$$E_t[a_{t+1}] = - \left[1 + \frac{1}{r(t+1)} \right] [(\bar{d}_t - \delta) - (\bar{d}_{M,t} - \delta_M) \beta]. \quad (3.34)$$

In general, this will deviate from zero. Similar to unexpected returns in the model with a single risky security, deviations from the CAPM are negatively related to past dividends and prices. In particular, for any asset i :

$$\begin{aligned} \text{cov}[p_{i,t}, a_{i,t+1}] &= - \left[1 + \frac{1}{r(t+1)} \right] \frac{1}{r t} [\text{var}(d_i) - \beta_i \text{cov}(d_i, d_M)] \\ &= - \left[1 + \frac{1}{r(t+1)} \right] \frac{1}{r t} \text{var}(\varepsilon_i), \end{aligned} \quad (3.35)$$

where $\text{var}(\varepsilon_i)$ is the residual variance when the asset's dividend is regressed on the market dividend. The ability of price to predict time-variation in $a_{i,t+1}$ is similar to its ability to predict unexpected returns, except that $\text{var}(\varepsilon_i)$ is substituted for the dividend's total variance. Again, we see that estimation risk induces price reversals and apparent overreaction by investors. When investors' best guess about expected dividends for a given stock is above the true mean (after adjusting for marketwide mispricing), price is inflated above its fundamental value and expected returns are lower than predicted by the CAPM.

Eq. (3.35) is essentially a time-series relation. The predictability of $a_{i,t+1}$ arises because investors do not know whether past dividends are greater than or less than the true mean. At any point in time, however, investors observe whether each security's average dividend is above or below the cross-sectional average. Our initial guess, then, was that deviations from the CAPM would not be *cross-sectionally* related to lagged prices: if cross-sectional variation in $a_{i,t+1}$ was related to the observable quantity $p_{i,t}$, it would seem that investors could use this information to earn abnormal returns. Surprisingly (to us), this intuition is wrong. In sample, the cross-sectional relation between $a_{i,t+1}$ and $p_{i,t}$ is

$$\text{cov}_{t+1}^{\text{cs}} [p_{i,t}, a_{i,t+1}] = \frac{1}{N} \sum_i (a_{i,t+1} - \bar{a}_{t+1}^{\text{cs}}) (p_{i,t} - \bar{p}_t^{\text{cs}}). \quad (3.36)$$

Taking the unconditional expectation yields

$$E[\text{cov}_{t+1}^{\text{cs}} (p_{i,t}, a_{i,t+1})] = \frac{1}{N} \sum_i \text{cov}(a_{i,t+1}, p_{i,t}) < 0, \quad (3.37)$$

which is negative because every covariance term is negative (see eq. 3.35). In the presence of estimation risk, lagged dividends and prices explain cross-sectional variation in expected returns

after controlling for betas. Investors understand the negative cross-sectional relation, but they cannot use this information to be better off.

We find this result paradoxical. To gain some intuition, consider the decision process of a rational investor. Implicitly, the expectation in eq. (3.37) integrates over all possible price paths from time 1 to $t+1$. However, at time t , the *conditional* cross-sectional relation can be either positive or negative, depending on the difference between \bar{d}_t and δ . In other words, conditional on observing \bar{d}_t , the cross-sectional covariance between prices and a_{t+1} depends on the true value of δ . Investors understand this dependence, and their beliefs about δ determine their investment choices. Thus, they integrate over the subjective distribution of δ to make portfolio decisions. The resulting belief about a_{t+1} will always have mean zero. The point is simply that investors do not ignore the relation between prices and deviations from the CAPM, but their best forecast of a_{t+1} at any point in time is always zero.

Alternatively, we can think about this in terms of an individual asset. Suppose that an asset has a relatively high price compared with other stocks. Does this imply that the asset is overvalued relative to its ‘fundamental’ value? The answer depends, of course, on the actual value of δ_i , which is unknown. Integrating over the posterior beliefs about δ_i , an investors’ best guess at all times is that the asset is fairly priced. Yet in hypothetical repeated sampling, the asset with the highest price will, on average, be overvalued. This puzzle highlights the distinction between the conditional nature of Bayesian decision making (conditional on the observed prices) and the frequentist perspective of classical statistics. For a Bayesian investor, hypothetical repeated sampling is irrelevant to the portfolio decision, which must be made after observing only a single realization of prices (see Berger, 1985, for an extensive discussion of these issues).

To illustrate the cross-sectional results, we simulate a set of prices and returns in the model. Similar to the example in Section 3.3, we assume that investors are risk-neutral and the

riskless rate is 0.05. In addition, all risky assets, with $N = 15$, have true expected dividends equal to 0.05. Hence, all prices equal one in the absence of estimation risk. When δ is unknown, security prices depend on realized dividends, which we randomly generate from a MVN distribution. To provide a reasonable covariance matrix, we estimate the return covariance matrix for 15 industry portfolios formed from all stocks on the Center for Research in Security Prices (CRSP) database.

Both the time-series and cross-sectional behavior of returns reveal the price-reversal effect of estimation risk. For $t = 10$ through 110, the correlation between total return and lagged price is negative for every security, with a mean correlation of -0.21. Deviations from the CAPM also appear predictable based on lagged prices: the average correlation between $a_{i,t+1}$ and $p_{i,t}$ is -0.16, and 14 out of the 15 correlations are negative. Cross-sectionally, the relation between $a_{i,t+1}$ and $p_{i,t}$ is significantly negative in Fama-MacBeth style regressions, with a t-statistic of -3.97. On average, an increase in price from one standard deviation below to one standard deviation above the cross-sectional mean leads to a -0.042 change in $a_{i,t+1}$. Since prices are generally close to one, this would imply that Jensen's alpha, based on rates of return, decreases by approximately -4.2%. Although investors attempt to hold mean-variance efficient portfolios and use all available information when making decisions, realized average returns can differ substantially from the predictions of the CAPM. Additional simulations show that this example is typical. For example, across 2500 simulations, Fama-MacBeth regressions produce an average t-statistic of -3.75 with a standard deviation of 0.94.

3.5. Informative priors, steady state, and simulations

We have presented an extremely simple model of estimation risk. Among the simplifications, we assumed that investors begin with no information about expected dividends, all parameters are constant, and dividends are IID. Each of these assumptions makes it difficult to judge the potential empirical significance of estimation risk. In this section, we relax the

assumptions to make the model a bit more realistic. We also present simulation evidence to suggest the practical importance of the results.

3.5.1. Informative priors

The assumption of diffuse priors has at least two important effects on the model. First, investors' beliefs about expected dividends are determined entirely by past realized dividends. With an informative prior, investors would put less weight on the data and more weight on their initial beliefs. Second, an investor's belief about the expected dividend on one asset is determined solely by the realized dividends on that asset, and does not depend at all on the realized payoffs of other securities. With an informative prior, however, dividends on assets with relatively high amounts of prior information can be useful in valuing other assets. We discuss both of these issues in this subsection. For now, we continue to assume that the true parameters of the dividend process remain constant over time.

Consider first the model with one risky asset. Assume that the variance of the dividend process, σ^2 , remains known, and suppose that investors begin with some information about the mean. In particular, assume that prior beliefs are centered around some δ^* and have variance σ^2/h , where h is a measure of prior information. Writing the variance in this form is simply for notational convenience; a variance equal to σ^2/h means that the investor has prior information that is as informative as a sample of h realized dividends. With this prior, a Bayesian investor's belief about dividends at time t is

$$d_{t+1} \sim^s N \left[\frac{h}{t+h} \delta^* + \frac{t}{t+h} \bar{d}_t, \frac{t+h+1}{t+h} \sigma^2 \right]. \quad (3.38)$$

Investors shrink their best guess about expected dividends toward their prior mean, and the variance reflects both the volatility of dividends, σ^2 , and uncertainty about the mean, $\sigma^2/(t+h)$. It is clear that the prior mean exerts a permanent, yet diminishing, influence on beliefs. To the extent that the prior mean deviates from δ , investors' beliefs are 'biased' away from the true

mean. However, as before, beliefs eventually converge to the true distribution as t gets large.

Equilibrium takes nearly the same form as the original model, except that price now reflects both the information in realized dividends as well as prior beliefs. Denote the mean of the subjective distribution as m_t . At time t , the price of the risky asset equals

$$p_t = \frac{1}{r} m_t - 2 \gamma f(t+h) \sigma^2$$

$$= \frac{1}{r} \frac{t}{t+h} \bar{d}_t + \frac{1}{r} \frac{h}{t+h} \delta^* - 2 \gamma f(t+h) \sigma^2, \quad (3.39)$$

where $f(t)$ is defined in eq. (3.11). With informative priors, the price contains a new term corresponding to the initial belief about expected dividends. It is clear from eq. (3.39) that the time-series properties of prices and returns will be determined by the properties of m_t . Moreover, the prior information anchors the price to the investor's initial guess, but does not have a stochastic effect on prices. As a result, in this simple model with fixed parameters, informative priors have little effect on the qualitative conclusions from the original model. Returns continue to be negatively related to past prices and dividends, although the magnitude is diminished compared with diffuse priors. For example,

$$\text{cov}[p_t, R_{t+1}] = -\frac{t}{r^2 (t+h)^2 (t+h+1)} \sigma^2, \quad (3.40)$$

which is negative but smaller than the corresponding expression with diffuse priors. This result is actually quite intuitive since prior information has, for practical purposes, the effect of simply adding h periods to the model before time 0.

We present more thorough simulations in Section 3.5.3, but it may be useful to report simulations here to illustrate the impact of informative priors. The model is simulated 2500 times assuming that investors are risk neutral, the riskless rate is 0.05, true expected dividends are 0.10, and the standard deviation of dividends is 0.10. Using the simulated data, we estimate the correlation between excess rates of return, which equal unexpected returns because of risk

neutrality, and lagged average dividends for 70 periods, from $t = 10$ through $t = 80$. The results confirm our analytic work. The average correlation equals -0.136 with perfect information (this is negative because of small-sample bias; see Stambaugh, 1999) and at the other extreme, the average correlation equals -0.259 with diffuse priors. Informative priors produce results that are between these polar cases. For example, with $h = 20$, meaning that investors have observed the equivalent of 30 periods of dividends when we begin estimating predictability, the average correlation between excess returns and lagged dividends equals -0.198 . These results are not sensitive to the prior mean.

We should add an important caveat at this point. The relatively minor effect of informative priors depends on the assumption that the true mean is fixed. Once we allow for shocks to the true parameters, informative priors can play a larger role because investors may appear to react slowly to changes in the dividend process (see the next section). In addition, notice that even in the current model, forecast errors are all expected to have the same sign because of the permanent influence of the prior mean. Although the influence is non-stochastic and does not affect serial correlation in returns, it could create the appearance of underreaction in some contexts. For example, Lewis (1989) argues that a similar phenomenon might account for the persistent forecast errors observed in the foreign exchange market in the 1980s.

Informative priors can also play a more important role with many assets. We need to consider two possible types of informative priors when there are many assets: symmetric information and differential information. In the discussion above, we denoted the variance of the prior as σ^2/h , where h can be interpreted as the length of the sample already observed. Loosely speaking, a symmetric prior means that the investor has observed the equivalent of h dividends for all securities. In this case, the prior covariance matrix equals Σ/h , where Σ is the covariance matrix of dividends. Of course, symmetric information is a fairly special case, and investors will typically have more information about some securities than others. With differential information, the prior covariance matrix does not have to be proportional to the

dividend covariance matrix.

We briefly consider the general case of differential information. Suppose that investors' prior beliefs about expected dividends are MVN $[\delta^*, \Omega]$. For a Bayesian investor, the posterior distribution for δ at time t is MVN $[m_t, \Pi_t]$, where

$$m_t = [\Omega^{-1} + t \Sigma^{-1}]^{-1} [\Omega^{-1} \delta^* + t \Sigma^{-1} \bar{d}_t], \quad (3.41)$$

$$\Pi_t = [\Omega^{-1} + t \Sigma^{-1}]^{-1}. \quad (3.42)$$

The mean of the distribution is a matrix-weighted-average of δ^* and \bar{d}_t , with weights given by the inverses of the covariance matrices. Importantly, the mean for a given asset will typically depend on the realized dividends for all assets and the covariance matrix does not have to be proportional to Σ . Beliefs about *dividends* (the previous equations are for δ) have the same mean; the covariance matrix reflects both the true variance of dividends and uncertainty about the mean, or $\Pi_t + \Sigma$.

In the special case of symmetric information, the mean and variance of the posterior distribution simplify to

$$m_t = \frac{h}{t+h} \delta^* + \frac{t}{t+h} \bar{d}_t, \quad (3.43)$$

$$\Pi_t = \frac{1}{t+h} \Sigma. \quad (3.44)$$

The mean is a scalar weighted-average of δ^* and \bar{d}_t , and an investor's belief about expected dividends for a given asset is unrelated to past dividends on other securities.

The equilibrium pricing function remains similar in form to the basic model. Specifically, the price at time t is

$$p_t = \frac{1}{r} m_t - 2 \gamma v(t) \iota, \quad (3.45)$$

where $v(t)$ is a deterministic $N \times N$ matrix that plays the role of $f(t)$ in the original model. Once

again, the properties of prices and returns depend on the behavior of m_t . Without going into too many details, we can draw two conclusions about the behavior of returns with informative priors:

(a) The cross-sectional correlation between deviations from the CAPM and lagged prices can be either positive or negative, depending on the strength of the prior and the relation between the prior mean and true expected dividends. Recall that with a diffuse prior, the cross-sectional correlation is always negative and investors appear to react too strongly to realized dividends. With an informative prior, however, investors can appear to update too slowly because they place less weight on the data and more on their prior beliefs.

To give a concrete example, suppose that the true mean of the dividend process is δ , an $N \times 1$ vector. Investors have symmetric priors and cannot distinguish among the assets, meaning that they have the same prior mean for every asset, or $\delta \sim^s N[\delta^* \mathbf{1}, \Sigma/h]$ where δ^* is a scalar and $\mathbf{1}$ is a vector of ones. To make matters simple, assume that the prior beliefs are correct on average, so that δ^* equals the cross-sectional average of δ . Under these assumptions, it can be shown that the expected cross-sectional covariance between deviations from the CAPM and lagged price equals

$$E[\text{cov}_{t+1}^{\text{cs}}(p_{i,t}, a_{i,t+1})] = \left[1 + \frac{1}{t+h+1}\right] \frac{t}{r(t+h)^2} [h \text{var}^{\text{cs}}(\delta) - \overline{\text{var}(\epsilon_i)}], \quad (3.46)$$

where $\text{var}(\epsilon_i)$ is the residual variance when the asset's dividend is regressed on the market dividend (see eq. 3.35), and $\overline{\text{var}(\epsilon_i)}$ denotes the cross-sectional average. The cross-sectional covariance can be either positive or negative depending on the strength of the prior (the parameter h) and the cross-sectional variance of δ . Qualitatively, these results are intuitive. When investors have weak prior beliefs (h is small), they appear to react too strongly to realized dividends and the price-reversal effect described in Section 3.3 dominates. On the other hand, with strong prior beliefs, investors rely less heavily on the data and might appear to react too

slowly to new information.

(b) In Section 3.4, we showed that estimation risk simply ‘scales up’ the return covariance matrix when investors have diffuse priors. This result does not have to hold with differential information. In the general model, the true conditional covariance matrix of returns is given by

$$\text{var}_t(\mathbf{R}_{t+1}) = \text{var}_t\left[\mathbf{d}_{t+1} + \frac{1}{r}\boldsymbol{\theta}_{t+1}\right] = \text{var}[\mathbf{M}_{t+1} \mathbf{d}_{t+1}] = \mathbf{M}_{t+1}\boldsymbol{\Sigma}\mathbf{M}'_{t+1}, \quad (3.47)$$

where

$$\mathbf{M}_{t+1} = \mathbf{I} + \frac{1}{r}\boldsymbol{\Pi}_{t+1}\boldsymbol{\Sigma}^{-1}. \quad (3.48)$$

The matrix \mathbf{M}_t maps unexpected dividends into unexpected returns. The identity matrix, \mathbf{I} , gives the immediate effect that unexpected dividends have on returns, and the second term gives the effect that unexpected dividends have on prices. The subjective covariance matrix of returns is:

$$\text{var}_t^s(\mathbf{R}_{t+1}) = \mathbf{M}_{t+1} [\boldsymbol{\Pi}_t + \boldsymbol{\Sigma}] \mathbf{M}_{t+1}^T. \quad (3.49)$$

The difference between the subjective and true covariance matrices is that the predictive covariance, $\boldsymbol{\Pi}_t + \boldsymbol{\Sigma}$, enters eq. (3.49).

Parameter uncertainty affects both the true and subjective distributions through the matrix \mathbf{M}_{t+1} . In general, the subjective and true covariance matrices will not be proportional to each other, nor will they be proportional to the covariance matrix when the dividend process is known. As a result, estimation risk affects subjective and true market betas differently, and both differ from market betas with perfect information. Consider, for example, a simple model with two assets, a low-information and a high-information security. Specifically, assume the investor has previously observed L periods of dividends for the low-information security and $H > L$ periods of dividends for the high-information security. In this case, it can be shown that parameter uncertainty increases the beta of the low-information security. Further, the subjective beta is greater than the true beta, implying that the true (observable) beta does not fully capture

the risk perceived by investors.

In summary, informative priors can be important for the way parameter uncertainty affects equilibrium prices and returns. Our basic conclusions about predictability and market efficiency, however, continue to hold.

3.5.2. Renewal of estimation risk

Perhaps the most obvious limitation of our model is that estimation risk steadily diminishes over time. As time passes, investors accumulate information and their beliefs converge to the true process. The reason is simple: we have assumed that the dividend process is fixed, so investors never ‘lose’ information. In reality, the economy evolves over time and a more realistic model would allow the dividend process to change. In this section, we extend the model to incorporate unobservable shocks to the true parameters which periodically renew estimation risk. We focus on the model with a single risky asset because the section is most applicable to the time-series properties of aggregate returns. At the microeconomic level, firms continually appear and disappear from the stock market, and it is not clear that the long-run implications of estimation risk are relevant for the behavior of individual stocks.

There are many ways to prevent estimation risk from vanishing in the limit. Here, we have chosen a particularly simple form of ‘renewal’ to illustrate the ideas. The model remains the same with one exception: we now assume that the true mean of the dividend process fluctuates over time at known, fixed intervals. Specifically, every K periods the mean is re-drawn from a normal distribution with mean δ^* and variance σ_s^2 . Thus, the model is essentially a sequence of short ‘regimes’ that look like our basic model truncated after K periods. We have analyzed alternative models in which (1) the length of the intervals is random rather than fixed and (2) the true mean of the dividend process follows a persistent process. The qualitative conclusions from these models appear to be similar.

After an infinite number of periods, it is clear that investors would learn the distribution from which the short-run mean is drawn. Therefore, in the limit, investors' priors at the beginning of each regime would be $N[\delta^*, \sigma_s^2]$. Although we analyze these priors as a special case, we do not think that it is either the most realistic or most interesting because it represents an extreme amount of learning. Instead, we consider the more general beliefs $N[\delta^*, \sigma_h^2]$, which have the same mean as the actual distribution but not necessarily the same variance. Thus, we assume that investors have observed the process long enough to know long-run expected dividends, even though they cannot observe short-run changes in the process. Permitting the variances to be different can be justified on several grounds.

First, we are trying to capture the idea that the economy moves through periods of high and low growth that cannot be perfectly observed. These periods might cover many years, so learning about the switching process – and its variance – is likely to be slow. Second, we have made the artificial assumption that the mean is repeatedly drawn from the same distribution. The economy undoubtedly moves through periods of relative stability and periods of rapid change, and the variance of shocks to expected dividends is likely to change over time. If investors cannot observe changes in volatility, then their current estimate of the volatility will not be perfect. Finally, alternative assumptions about the evolution of the true mean do not necessarily have the property that the prior variance ever converges to the true variance.²³ We abstract from these issues, and take the more expeditious approach of simply permitting the prior variance to be different from σ_s^2 .

The pricing function is similar to the price in the basic model. The renewal model consists of a sequence of intervals with fixed expected dividends, and investors do not observe the current draw of the short-run mean, δ_k . As discussed above, investors' priors at the beginning

²³ For example, suppose the dividend mean $\hat{\delta}$ follows a random walk, dividends have conditional variance σ^2 , and the shocks to $\hat{\delta}$ are uncorrelated with dividends and have variance σ_s^2 . In the long-run, investors beliefs about $\hat{\delta}$ will be $N[m, \sigma_h^2]$, where σ_h^2 is time-invariant and $\sigma_h^2 > \sigma_s^2$.

of each interval are $N[\delta^*, \sigma_h^2]$. For notational convenience, let $\sigma_s^2 = \sigma^2/s$ and $\sigma_h^2 = \sigma^2/h$, and assume for simplicity that investors are risk neutral. Realized dividends during the current interval provide no information about payoffs after the end of the interval, so beliefs about those payoffs always have mean δ^* . Therefore, the price at the beginning of every regime equals δ^*/r , the value of expected dividends in perpetuity. After t periods in the current regime, the investor's predictive belief about short-run dividends has mean

$$m_t = \frac{h}{t+h} \delta^* + \frac{t}{t+h} \bar{d}_t, \quad (3.50)$$

identical to eq. (3.38). Thus, price equals

$$p_t = AF_{K-t} m_t + \frac{1}{(1+r)^{K-t}} \frac{\delta^*}{r}, \quad (3.51)$$

where AF_{K-t} is an annuity factor for $K-t$ periods. Not surprisingly, the time-series properties of prices and returns once again depend on the behavior of m_t . It is straightforward to show that excess, or unexpected, returns are given by

$$UR_{t+1} = \left[1 + \frac{AF_{K-t-1}}{t+h+1} \right] (d_{t+1} - m_t). \quad (3.52)$$

The term in parentheses is simply unexpected dividends, which have an immediate effect on unexpected returns (the '1' in brackets) and an indirect effect on prices (with the multiplier $AF_{K-t-1}/(t+h+1)$).

The analysis of predictability with renewal is more complicated than in our basic model. In particular, now that the short-run mean is random, we have to distinguish between expectations that are conditional on the current mean and expectations that treat the parameter as random. It turns out that a combination of the two seems to be relevant for empirical tests (see the simulations below). At time t (interpreted as t periods into the current regime), the unexpected return has true mean

$$E_t[UR_{t+1}] = \left[1 + \frac{AF_{K-t-1}}{t+h+1} \right] (\delta_k - m_t), \quad (3.53)$$

which follows immediately from eq. (3.52). As in our basic model, the true unexpected return is negatively related to past dividends and prices. Consequently, *taking the value of d_k as given*, the covariance between excess returns and lagged prices equals

$$\begin{aligned} \text{cov}[p_t, UR_{t+1}] &= -AF_{K-t} \left(1 + \frac{AF_{K-t-1}}{t+h+1} \right) \text{var}(m_t) \\ &= -AF_{K-t} \left(1 + \frac{AF_{K-t-1}}{t+h+1} \right) \left(\frac{t}{t+h} \right)^2 \frac{1}{t} \sigma^2, \end{aligned} \quad (3.54)$$

which is negative. We refer to this expression as the ‘conditional covariance’ because it regards the short-run mean as fixed. The equation is very similar to our previous results with informative priors, except that the covariance is attenuated because price fluctuations are less pronounced (the price always returns at the end of the regime to δ^*/r). Therefore, in one sense, the effects of estimation risk documented above remain the same even in the long-run: the true and subjective distributions are different, leading to price reversals.

Unfortunately, things are not quite so simple. Although the conditional covariance does not depend on δ_k , the ‘unconditional covariance’ – which regards the short-run mean as random – will nonetheless differ from eq. (3.54).²⁴ Specifically, the unconditional covariance equals

$$\begin{aligned} \text{cov}[p_t, UR_{t+1}] &= AF_{K-t} \left(1 + \frac{AF_{K-t-1}}{t+h+1} \right) [\text{cov}(d_{t+1}, m_t) - \text{var}(m_t)], \\ &= AF_{K-t} \left(1 + \frac{AF_{K-t-1}}{t+h+1} \right) \frac{t}{(t+h)s} \left[1 - \frac{t+s}{t+h} \right] \sigma^2. \end{aligned} \quad (3.55)$$

The sign of the unconditional covariance depends on the relative magnitudes of s and h . Recall that σ^2/s is the true variance of δ_k while σ^2/h is the prior variance. Therefore, the unconditional covariance is negative when the prior variance is greater than the true ($h < s$), but positive when

²⁴ In statistical terms, the expected conditional covariance does not equal the unconditional covariance because the means of the variables move together over time.

the prior variance is less. When investors believe that the variance of shocks to expected dividends is high, they are relatively sensitive to realized dividends and the price-reversal effect of estimation risk shows up both conditionally and unconditionally. On the other hand, if the short-run mean is more variable than investors believe, they tend to be surprised by the large movements in expected dividends and require many observations to update their beliefs. Consequently, returns exhibit patterns of continuation or momentum. The cutoff value occurs when investors have exactly the right beliefs about the variance of δ_k , or when $s = h$. In this case, the unconditional covariance between excess returns and lagged prices is exactly zero.

Thus, we have two results on predictability in the renewal model: (1) the conditional covariance is always negative, regardless of the relative magnitudes of s and h , and (2) the unconditional covariance depends on whether h is less than or greater than s . The fact that the conditional covariance is negative implies immediately that excess returns are predictable, but it is not obvious to us whether the unconditional or conditional covariance is more relevant for standard empirical tests.²⁵ An empirical test depends on the observed sample, and implicitly conditions on the sample value (or values) of the mean parameter δ_k . This observation suggests that the conditional variance might be most relevant. Indeed, take a particularly simple case in which observed sample covers only one regime. Regardless of the value of δ_k , the covariance between unexpected returns and prices is expected to be negative; the correlation in this case corresponds directly to the conditional covariance.²⁶ On the other hand, if a sample covers multiple regimes, the empiricist implicitly conditions on several values of δ_k and our simple formula for the conditional variance no longer represents the population counterpart of the estimate. To muddy the waters further, if the empiricist suspects that a change in regime occurs

²⁵ Some additional explanation might be useful. A predictive regression for returns that includes regime dummies would estimate the conditional covariance, and can therefore detect the price reversals. Alternatively, the price reversals can be picked up by estimating within-regime covariances. However, it is not common to include regime dummies in predictive regressions, nor is it easy to identify regime changes.

²⁶ We stress that this is not a survival bias or a so-called ‘peso problem.’ We expect to see a negative correlation *for any value of d_k* because the true correlation is negative.

and adds a dummy variable to the regression, or focuses on subperiod regressions, then the sample covariance will correspond once again to the conditional variance. Rather than speculate further, we rely on simulations to assess the expected value of the sample covariance with estimation risk.

3.5.3. Simulations

To investigate the ‘steady-state’ effects of estimation risk with renewal, we simulate the model 2500 times and examine the predictability of returns. To make the model more realistic, the simulations assume that dividends follow a geometric random walk with time-varying growth. Specifically, dividends follow the process

$$\ln d_{t+1} = g_k + \ln d_t + \varepsilon_{t+1}, \quad (3.56)$$

where $\varepsilon_{t+1} \sim N[0, \sigma^2]$ and g_k is randomly drawn every K periods from a normal distribution with mean g^* and variance σ^2/s . The simulations normalize the initial dividend to equal one, the discount rate equals 0.12, $\sigma = 0.10$, and the long-run growth rate g^* equals 0.03. These parameters are chosen to be reasonably close to actual values, interpreting a period in the model as one year. In comparison, the average annual return on the CRSP value-weighted index equals 12.5% for the period 1926 through 1997, and Brennan and Xia (1998) report that the average real growth rate in dividends equals 1.6%, with a standard deviation of 12.9%, over the period 1871-1996. The simulations estimate predictive regressions using roughly 75 years of data, again taken to be similar to a typical study. We report results for several combinations of the parameters s , h , and K . These parameters determine the true variance in short-run growth rates, the variance of investors’ priors, and the length of a regime, respectively. Appendix B describes the Bayesian inference problem for this model.

Table 3.1 reports the results of the simulations. Specifically, the table shows the average slope coefficient and t-statistic when excess returns are regressed on lagged dividend yield. An

important complication arises because the slope coefficient in these regressions suffers from a significant small-sample bias (see Stambaugh, 1999, and the discussion in Section 2.3). The bias is caused by the same phenomenon that biases autocorrelation estimates downward, but the coefficients in these regressions are biased upward, giving the appearance of more predictability. To help reduce the effects of the bias, we also report bias-adjusted slope coefficients using the results of Stambaugh (1999).²⁷ In addition, the table reports results when investors perfectly observe the dividend process. The difference between the bias-adjusted coefficients with estimation risk and with perfect information represents an estimate of the predictability caused by estimation risk.

Table 3.1 shows that estimation risk can induce predictability even in steady state. The results suggest that the negative conditional covariance tends to dominate the regressions. Even when investors know both the mean and variance of the distribution from which the growth rate is drawn ($h = s$), the slope coefficient in the dividend yield regression is positive. For example, with two regimes over the 75 years, the average slope coefficient ranges from 1.06 to 1.52 for different values of $h = s$ (see the diagonal terms in the last column). With four regimes the slope coefficient ranges from 0.56 to 0.60, and with six regimes the slope ranges from 0.45 to 0.50. The price reversal effect tends to be larger when the regimes are longer, and it becomes much more pronounced when investors' prior variance is higher than actual variance. With $s = 49$ and $h = 16$, the table shows that the slope coefficient varies between 1.98 and 2.24 for different values of K . Cases in which $s > h$, so the subjective variance is greater than the true, is of particular interest because it shows roughly how prices behave before we reach steady state

²⁷ To derive the bias, Stambaugh (1999) makes several assumptions about the return and dividend processes that do not hold in our model (e.g., dividend yields are AR(1) and returns are homoskedastic). Indeed, Table 1 shows that with perfect information, the bias adjustment tends to correct too much (the corrected slopes are negative not zero). To confirm that the simulation evidence is not driven by problems with the bias-adjustment procedure, we perform an additional check. We also estimate regressions using true unexpected returns, which always have conditional mean zero but otherwise have the same properties as excess returns. The average slope coefficient in these regressions provides an alternative estimate of the bias. These results support our conclusions in Table 3.1.

Table 3.1
Predictability in steady state

We simulate the renewal model 2500 times. Dividends are assumed to follow a geometric random walk with time-varying expected growth, where the short-run growth rate g_k is randomly drawn every K periods from $N[g^*, \sigma^2/s]$. Investors are risk neutral and have initial beliefs about g_k at the beginning of each regime equal to $N[g^*, \sigma^2/h]$. In the simulations, $r = 0.12$, $\sigma = 0.10$, and $g^* = 0.03$. The table reports various combinations of s , h , and K . The table shows the average slope coefficient and t -statistic when excess returns are regressed on lagged dividend yield for roughly 75 years (we require the number of years to be divisible by K). We also report bias-adjusted slope coefficients which correct for small-sample bias using the results of Stambaugh (1999).

		Estimation risk			Perfect information			Difference				
		h			h			h				
		s	16	25	49	16	25	49	16	25	49	
2 regimes (K = 38)	slope	16	2.91	3.05	3.68	0.40	0.39	0.50	2.41	2.65	3.16	
		25	3.31	3.57	4.41	0.37	0.31	0.52	2.94	3.26	3.88	
		49	3.92	4.16	5.62	0.44	0.50	0.41	3.48	3.65	5.21	
	bias-adj slope	16	1.6	0.78	0.32	-0.53	-0.28	-0.29	-0.16	1.06	0.60	-0.37
		25	1.6	1.16	0.84	0.15	-0.28	-0.34	-0.11	1.44	1.18	0.25
		49	1.6	1.77	1.39	1.31	-0.21	-0.14	-0.21	1.98	1.53	1.52
	t-stat	16	1.02	0.80	0.53	0.25	0.25	0.28	0.77	0.55	0.25	
		25	1.15	0.98	0.71	0.23	0.19	0.18	0.92	0.78	0.53	
		49	1.34	1.14	0.94	0.11	0.14	0.13	1.23	1.00	0.81	
4 regimes (K = 19)	slope	16	2.56	2.44	1.97	0.37	0.30	0.34	2.19	2.15	1.64	
		25	3.34	3.29	3.61	0.32	0.32	0.28	3.02	2.97	3.32	
		49	4.01	4.21	5.25	0.35	0.35	0.27	3.65	3.85	4.98	
	bias-adj slope	16	1.6	0.29	-0.56	-2.95	-0.27	-0.33	-0.29	0.56	-0.23	-2.66
		25	1.6	1.06	0.26	-1.36	-0.29	-0.31	-0.34	1.35	0.57	-1.01
		49	1.6	1.72	1.17	0.24	-0.27	-0.27	-0.35	1.99	1.43	0.60
	t-stat	16	0.62	0.38	0.05	0.24	0.20	0.22	0.38	0.18	-0.17	
		25	0.86	0.59	0.32	0.18	0.18	0.16	0.69	0.41	0.16	
		49	1.05	0.80	0.56	0.14	0.14	0.11	0.91	0.66	0.46	
6 regimes (K = 13)	slope	16	2.59	1.94	1.41	0.27	0.32	0.24	2.33	1.62	1.17	
		25	3.65	3.44	3.65	0.31	0.28	0.25	3.34	3.16	3.40	
		49	4.38	4.87	5.80	0.29	0.30	0.26	4.09	4.57	5.54	
	bias-adj slope	16	1.6	0.16	-1.35	-4.15	-0.33	-0.26	-0.34	0.49	-1.08	-3.81
		25	1.6	1.22	0.14	-1.95	-0.27	-0.31	-0.34	1.49	0.45	-1.61
		49	1.6	1.94	1.56	0.18	-0.29	-0.27	-0.32	2.24	1.83	0.50
	t-stat	16	0.48	0.19	-0.05	0.17	0.20	0.16	0.31	-0.01	-0.20	
		25	0.42	0.45	0.22	0.16	0.15	0.13	0.57	0.31	0.08	
		49	0.90	0.71	0.44	0.11	0.10	0.10	0.79	0.60	0.34	

(even if investors know σ_s^2 , the subjective variance of dividends is always greater than the true after a finite number of periods). We believe that the evolutionary process is as relevant for empirical tests as the steady-state equilibrium.

To add some perspective, the historical slope coefficient for the period 1941 to 1997 is 3.93 (standard error of 1.73), before adjusting for bias, when the CRSP value-weighted return is regressed on its lagged dividend yield. Although a more thorough study is necessary to draw detailed conclusions, the simulations provide preliminary evidence that estimation risk could account for a non-trivial portion of the predictability. We hesitate to draw firm conclusions because the simulations do not (and probably cannot) capture all of the relevant properties of actual dividends and returns, and it is beyond the scope of the current paper to understand which set of parameter values best characterizes the historical stock market.

The table also shows that return continuation, or a negative slope coefficient in the dividend yield regressions, is possible if investors' prior variance is smaller than the true. This case corresponds to a situation in which the economy is changing more dramatically than investors anticipate. Investors require many dividend observations until their beliefs 'catch up' with the actual changes, which creates persistence in expected returns.

Finally, adding a regime dummy variable to the regressions produces an estimate of the conditional covariance. In results not reported, the average bias-adjusted slope coefficient is approximately 1.47 with two regimes, 2.00 with four regimes, and 2.55 with six regimes. These values are not sensitive to the values of h and s , presumably because h and s affect the covariance in the numerator and the variance in the denominator by similar magnitudes. Although we believe these issues deserve a more complete treatment, we simply note here that the simulations confirm, in substance, our earlier results. Even in steady state, parameter uncertainty can be a source of predictability.

3.6. Summary and conclusions

Financial economists generally assume that, unlike themselves, investors know the means, variances, and covariances of the return or cashflow process. Practitioners do not have this luxury. To apply the elegant framework of modern portfolio theory, they must estimate expected returns using whatever information is available. As Black (1986) observes, however, the world is a noisy place and our observations are necessarily imprecise. The estimation risk literature formalizes this problem. Surprisingly, this literature has had little impact on mainstream thinking about equilibrium asset pricing and market efficiency. We believe that this is due, in large part, to its focus on the subjective beliefs of investors, rather than the true, or empirical, distribution of returns. As we have emphasized throughout the paper, the subjective distribution of returns does not have to correspond to the empirical distribution even when investors are rational.

Our analysis shows that parameter uncertainty can significantly affect the time-series and cross-sectional behavior of asset prices. Prices in our model satisfy commonly accepted notions of market efficiency and rational expectations: investors use all available information when making decisions and, in equilibrium, the perceived pricing function equals the true pricing function. However, prices and returns violate standard tests of efficiency, suggesting that parameter uncertainty is likely to be important for characterizing an efficient market. Although we do not argue that estimation risk necessarily explains specific asset-pricing anomalies, our results relate to several empirically-observed patterns in stock prices:

Return predictability. Empirical studies document time-varying expected stock returns, captured by variables like past returns, aggregate dividend yield, and aggregate book-to-market (e.g., Keim and Stambaugh, 1986; Fama and French, 1989; Kothari and Shanken, 1997). These studies attribute variation in expected returns to changes in business conditions or to irrational investors. We find that estimation risk can be a third source of return predictability. In our basic model, expected returns are negatively related to past prices, dividends, and returns. The

price-reversal effects become more pronounced in long-horizon returns, consistent with the evidence of Fama and French (1988) and Poterba and Summers (1988). In more elaborate models, parameter uncertainty could also give the appearance of underreaction or momentum in returns.

Volatility. Leroy and Porter (1981) and Shiller (1981) derive bounds on the volatility of asset prices in an efficient market. They conclude that prices ‘move too much to be justified by subsequent changes in dividends.’ Our findings suggest that estimation risk might help explain excess volatility. Asset prices can reject the volatility bounds even though investors are rational and prices reflect all available information. The volatility bounds can be viewed as tests of market efficiency only if investors have perfect knowledge of the dividend process. In our simple model with IID dividends, price changes are completely uncorrelated with future dividends. Thus, like the results on predictability, price volatility would suggest investor overreaction in the absence of estimation risk. Asset prices can take long swings away from ‘fundamental’ value, which are eventually reversed, giving the appearance of fads or bubbles in stock prices.

CAPM. Many empirical studies find that the CAPM does not completely describe the cross-section of expected returns. Departures from the CAPM have been attributed to missing risk factors, irrational investors, or trading frictions. We find that estimation risk provides an additional explanation. When investors must estimate expected dividends, returns will typically deviate from the predictions of the CAPM even if investors attempt to hold mean-variance efficient portfolios. Moreover, the deviations can be predictable, both cross-sectionally and in time series, with past dividends, prices, and returns. Our results complement previous studies on asset pricing with incomplete information (e.g., Williams, 1977).

The fact that estimation risk might explain these patterns does not, of course, mean that it does. The impact of estimation risk on actual prices is obviously an empirical issue, which we plan to explore in future work. Clarkson and Thompson (1990) find evidence that market betas

reflect differences in the quality of available information about firms, consistent with differentially-informative priors. However, our analysis suggests the possibility of much more general effects on volatility and predictability, at both the individual-security and aggregate-market levels. The central question becomes: To what extent do rational forecasts deviate from expectations based on perfect knowledge of the underlying cashflow process? We believe, from casual observation and reading of the financial press, that these deviations could be quite large. To assess market efficiency in light of estimation risk, the researcher may, in effect, need to mimic the Bayesian-updating process of rational investors.

It is important to distinguish between ‘true’ uncertainty in the economy and estimation risk. True uncertainty concerns economic conditions or events that could not be predicted even with complete knowledge of the underlying economic process. In contrast, estimation risk refers to subjective uncertainty about some relevant characteristic of the economy that is already largely determined at the time of the forecast, but not directly observable. Although the line between subjective and true is not always clear, the distinction can be important for asset pricing. As we have seen, uncertainty about a predetermined characteristic (expected dividends in our model) gives rise to price-related predictability in returns, since resolution of this uncertainty is negatively related to past mistakes. In contrast, resolution of true uncertainty will be unrelated to past information.

As an example of subjective uncertainty, consider the rate of productivity growth in the United States, which has recently received much attention. Market analysts debate whether past technological innovations allow the economy to grow more quickly. The question, then, is whether productivity growth has already accelerated; the change in the economy is presumed to have already taken place, but it is unknown. Similarly, Lewis (1989) argues that demand for U.S. currency shifted in the early 1980s, but investors could not immediately learn about this change. At the firm level, uncertainty about the demand for a firm’s product or service would generate estimation risk. Consumers’ preferences, and consequently true expected demand,

might be predetermined, but ‘noise’ prevents investors from precisely measuring the true probability distribution of demand. In all of these examples, the underlying economic process cannot be perfectly observed.

We close with a few reflections on the relation between data mining and estimation risk. In recent years, researchers and practitioners have become increasingly sensitive to the possibility that, with the intensive scrutiny of data common in investment research, ‘statistically significant’ return patterns can emerge even when returns are essentially random (see, for example, Lo and MacKinlay, 1990). Thus, we might observe patterns that do not exist in the true underlying process. Our analysis of estimation risk suggests a complementary concern. With hindsight, we can discern patterns that existed in the true return process, but could not have been exploited at the time by rational investors. Similar to the results of data snooping, these patterns would not be relevant for future investment decisions. Unlike data snooping, however, the patterns can persist in the future because they are part of the true process. This conclusion provides an alternative perspective on empirical anomalies. For example, Fama (1998) argues that various long-horizon return anomalies in the literature are chance results, consistent with market efficiency. He finds that ‘apparent overreaction to information is about as common as underreaction’ and, given data mining and other methodological concerns, concludes that the overall weight of the evidence is not compelling. Our work reinforces this conclusion by demonstrating that reversals and continuations might be expected in an efficient market with estimation risk, not only as a random outcome of the data but as a feature of the actual process.

On the Predictability of Stock Returns: Theory and Evidence

References

- Abel, Andrew and Frederic Mishkin, 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.
- Ball, Ray, 1978, Anomalies in relationships between securities' yields and yield-surrogates, *Journal of Financial Economics* 6, 103-126.
- Barberis, Nicholas, 2000, Investing for the long run when returns are predictable, *Journal of Finance* 55, 225-264.
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 307-343.
- Bawa, Vijay and Stephen Brown, 1979, Capital market equilibrium: Does estimation risk really matter?, in: V. Bawa, S. Brown, and R. Klein, eds., *Estimation Risk and Optimal Portfolio Choice* (North-Holland, Amsterdam).
- Bawa, Vijay, Stephen Brown, and Roger Klein, 1979, *Estimation Risk and Optimal Portfolio Choice* (North-Holland, Amsterdam).
- Berger, James, 1985, *Statistical Decision Theory and Bayesian Analysis* (Springer-Verlag, New York, NY).
- Berk, Jonathan, 1995, A critique of size-related anomalies, *Review of Financial Studies* 8, 275-286.
- Black, Fischer, 1972, Capital market equilibrium with restricted borrowing, *Journal of Business* 45, 444-455.
- Black, Fischer, 1986, Noise, *Journal of Finance* 41, 529-543.
- Black, Fischer, Michael Jensen, and Myron Scholes, 1972, The capital asset pricing model: Some empirical tests, in: M. Jensen, ed., *Studies in the theory of capital markets* (Praeger, New York, NY), 79-121.
- Bossaerts, Peter, 1997, The dynamics of equity prices in fallible markets, Working paper (California Institute of Technology, Pasadena, California).
- Breeden, Douglas, 1979, An intertemporal asset pricing model with stochastic consumption and investment opportunities, *Journal of Financial Economics* 7, 265-296.
- Brennan, Michael and Yihong Xia, 1998, Stock price volatility, learning, and the equity premium, Working paper (University of California at Los Angeles, Los Angeles, CA).
- Campbell, John, 1991, A variance decomposition for stock returns, *The Economic Journal* 101, 157-179.

- Chan, K.C. and Nai-fu Chen, 1991, Structural and return characteristics of small and large firms, *Journal of Finance* 46, 1467-1484.
- Chan, Louis, Yasushi Hamao, and Josef Lakonishok, 1991, Fundamentals and stock returns in Japan, *Journal of Finance* 43, 309-325.
- Clarkson, Peter, Jose Guedes, and Rex Thompson, 1996, On the diversification, observability, and measurement of estimation risk, *Journal of Financial and Quantitative Analysis* 31, 69-84.
- Clarkson, Peter and Rex Thompson, 1990, Empirical estimates of beta when investors face estimation risk, *Journal of Finance* 45, 431-453.
- Coles, Jeffrey and Uri Loewenstein, 1988, Equilibrium pricing and portfolio composition in the presence of uncertain parameters, *Journal of Financial Economics* 22, 279-303.
- Coles, Jeffrey, Uri Loewenstein, and Jose Suay, 1995, On equilibrium pricing under parameter uncertainty, *Journal of Financial and Quantitative Analysis* 30, 347-364.
- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and over-reactions, *Journal of Finance* 53, 1839-1885.
- Daniel, Kent and Sheridan Titman, 1997, Evidence on the characteristics of cross-sectional variation in stock returns, *Journal of Finance* 52, 1-33.
- Davis, James, 1994, The cross-section of realized stock returns: The pre-Compustat evidence, *Journal of Finance* 50, 1579-1593.
- DeLong, J. Bradford, Andrei Shleifer, Lawrence Summers, and Robert Waldmann, 1990, Noise trader risk in financial markets, *Journal of Political Economy* 98, 703-738.
- Detemple, Jerome, 1986, Asset pricing in a production economy with incomplete information, *Journal of Finance* 41, 383-391.
- Dothan, Michael and David Feldman, 1986, Equilibrium interest rates and multiperiod bonds in a partially observable economy, *Journal of Finance* 41, 369-382.
- Fama, Eugene, 1976, *Foundations of Finance* (Basic Books, New York, NY).
- Fama, Eugene, 1991, Efficient capital markets: II, *Journal of Finance* 46, 1575-1617.
- Fama, Eugene and Kenneth French, 1988, Permanent and temporary components of stock prices, *Journal of Political Economy* 96, 246-273.
- Fama, Eugene and Kenneth French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23-49.
- Fama, Eugene and Kenneth French, 1992, The cross-section of expected stock returns, *Journal of Finance* 47, 427-465.

- Fama, Eugene and Kenneth French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3-56.
- Fama, Eugene and Kenneth French, 1997, Industry costs of equity, *Journal of Financial Economics* 43, 153-193.
- Fama, Eugene and James MacBeth, 1973, Risk, return and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607-636.
- Fama, Eugene and G. William Schwert, 1977, Asset returns and inflation, *Journal of Financial Economics* 5, 115-146.
- Genotte, Gerard, 1986, Optimal portfolio choice under incomplete information, *Journal of Finance* 41, 733-746.
- Gibbons, Michael, Stephen Ross, and Jay Shanken, 1989, A test of the efficiency of a given portfolio, *Econometrica* 57, 1121-1152.
- Greene, William, 1993, *Econometric Analysis* (Macmillan Publishing, New York, NY).
- Harvey, Campbell, 1989, Time-varying conditional covariances in tests of asset pricing models, *Journal of Financial Economics* 24, 289-317.
- Haugen, Robert and Nardin Baker, 1996, Commonality in the determinants of expected stock returns, *Journal of Financial Economics* 41, 401-440.
- He, Jia, Raymond Kan, Lilian Ng, and Chu Zhang, 1996, Tests of the relations among marketwide factors, firm-specific variables, and stock returns using a conditional asset pricing model, *Journal of Finance* 51, 1981-1908.
- Jegadeesh, Narasimhan and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65-91.
- Jobson, J.D., Bob Korkie, and V. Ratti, 1979, Improved estimation for Markowitz portfolios using James-Stein type estimators, *Proceedings of the American Statistical Association*, 279-284.
- Johnson, Richard and Dean Wichern, 1982, *Applied Multivariate Statistical Analysis* (Prentice Hall, Englewood Cliffs, NJ).
- Johnston, John, 1984, *Econometric Methods* (McGraw-Hill, New York, NY).
- Jorion, Philippe, 1985, International portfolio decisions with estimation risk, *Journal of Business* 58, 259-278.
- Kandel, Shmuel and Robert Stambaugh, 1996, On the predictability of stock returns: An asset allocation perspective, *Journal of Finance* 51, 385-424.
- Keim, Donald and Robert Stambaugh, 1986, Predicting returns in the stock and bond markets, *Journal of Financial Economics* 17, 357-390.

- Kothari, S.P. and Jay Shanken, 1997, Book-to-market, dividend yield, and expected market returns: A time-series analysis, *Journal of Financial Economics* 44, 169-203.
- Kothari, S.P., Jay Shanken, and Richard Sloan, 1995, Another look at the cross-section of expected stock returns, *Journal of Finance* 50, 185-224.
- Lakonishok, Josef, Andrei Shleifer, and Robert Vishny, 1994, Contrarian investment, extrapolation, and risk, *Journal of Finance* 49, 1541-1578.
- LeRoy, Stephen, 1973, Risk aversion and the martingale property of stock prices, *International Economic Review* 14, 436-446.
- LeRoy, Stephen and Richard Porter, 1981, The present value relation: Tests based on implied variance bounds, *Econometrica* 49, 555-574.
- Lewis, Karen, 1989, Changing beliefs and systematic rational forecast errors with evidence from foreign exchange, *American Economic Review* 79, 621-636.
- Lintner, John, 1965, The valuation of risky assets and the selection of risky investments in stock portfolios and capital budgets, *Review of Economics and Statistics* 47, 13-37.
- Lo, Andrew and A. Craig MacKinlay, 1990, Data-snooping biases in tests of financial asset pricing models, *Review of Financial Studies* 3, 431-467.
- Lucas, Robert, 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429-1446.
- MacKinlay, A. Craig, 1995, Multifactor models do not explain deviations from the CAPM, *Journal of Financial Economics* 38, 3-28.
- McEnally, Richard and Rebecca Todd, 1993, Systematic risk behavior of financially distressed firms, *Quarterly Journal of Business and Economics* 32, 3-19.
- Merton, Robert, 1971, Optimum consumption and portfolio rules in a continuous-time model, *Journal of Economic Theory* 3, 373-413.
- Merton, Robert, 1973, An intertemporal asset pricing model, *Econometrica* 41, 867-887.
- Muth, John, 1961, Rational expectations and the theory of price movements, *Econometrica* 29, 315-335.
- Pontiff, Jeffrey and Lawrence Schall, 1998, Book-to-market ratios as predictors of market returns, *Journal of Financial Economics* 49, 141-160.
- Poterba, James and Lawrence Summers, 1988, Mean reversion in stock prices: Evidence and implications, *Journal of Financial Economics* 22, 27-59.
- Roll, Richard, 1977, A critique of the asset pricing theory's tests – Part 1: On past and potential testability of the theory, *Journal of Financial Economics* 4, 129-176.

- Rosenberg, Barr, Kenneth Reid, and Ronald Lanstein, 1985, Persuasive evidence of market inefficiency, *Journal of Portfolio Management* 11, 9-17.
- Shanken, Jay, 1987, Multivariate proxies and asset pricing relations: Living with the Roll critique, *Journal of Financial Economics* 18, 91-110.
- Shanken, Jay, 1990, Intertemporal asset pricing: An empirical investigation, *Journal of Econometrics* 45, 99-120.
- Sharpe, William F., 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, *Journal of Finance* 19, 425-442.
- Shiller, Robert, 1981, Do stock prices move too much to be justified by subsequent changes in dividends?, *American Economic Review* 71, 421-436.
- Stambaugh, Robert, 1997, Analyzing investments whose histories differ in length, *Journal of Financial Economics* 45, 285-331.
- Stambaugh, Robert, 1999, Predictive regressions, *Journal of Financial Economics* 54, 375-421.
- Stattman, Dennis, 1980, Book values and stock returns, *The Chicago MBA: A Journal of Selected Papers* 4, 25-45.
- Stulz, René, 1987, An equilibrium model of exchange rate determination and asset pricing with nontraded goods and imperfect information, *Journal of Political Economy* 95, 1024-1040.
- Timmermann, Allan, 1993, How learning in financial markets generates excess volatility and predictability in stock prices, *Quarterly Journal of Economics* 108, 1135-1145.
- Timmermann, Allan, 1996, Excess volatility and predictability of stock prices in autoregressive dividend models with learning, *Review of Economic Studies* 63, 523-557.
- Wang, Jiang, 1993, A model of intertemporal asset prices under asymmetric information, *Review of Economic Studies*, 60, 249-282.
- White, Halbert, 1984, *Asymptotic theory for econometricians* (Academic Press, Orlando, FL).
- Williams, Joseph, 1977, Capital asset prices with heterogeneous beliefs, *Journal of Financial Economics* 5, 219-239.
- Zellner, Arnold, 1962, An efficient method of estimating seemingly unrelated regressions and tests of aggregation bias, *Journal of the American Statistical Association* 57, 500-509.
- Zellner, Arnold, 1971, *An Introduction to Bayesian Inference in Econometrics* (John Wiley and Sons, New York, NY).

On the Predictability of Stock Returns: Theory and Evidence

Appendix A

This appendix supplements Chapter 2. I prove that h_i equals zero in eq. (2.5) and give a specific example in which asset prices satisfy the proposition. In addition, the appendix describes the Fama and French (1993) factors used in the empirical tests and summarizes the bootstrap simulations in Section 2.3.

A.1. Proof that $h_i = 0$

Let M be the proxy for the market portfolio, and assume that HML is constructed so that M and HML span the conditional tangency portfolio. The portfolio weights of HML can change over time, but I suppress the time subscript for simplicity. Without lack of generality, assume that $\text{cov}(R_M, \text{HML}) = 0$. We need to show that, under the mispricing story, the factor loading on HML must be zero in the unconditional time-series regression

$$R_i(t) = a_i + b_i R_M(t) + h_i \text{HML}(t) + \epsilon_i(t). \quad (\text{A.1})$$

I assume that mispricing is temporary, by which I mean that conditional deviations from the CAPM have expectation zero. Also, assume that any time-variation in the conditional factor loadings is unrelated to time-variation in the factors' expected returns. These assumptions imply that the CAPM holds unconditionally:

$$E[R_i] = b_i' E[R_M], \quad (\text{A.2})$$

where b_i' is the unconditional market beta. Also, taking expectations in eq. (A.1) yields:

$$E[R_i] = a_i + b_i E[R_M] + h_i E[\text{HML}]. \quad (\text{A.3})$$

If $a_i = 0$ and $b_i = b_i'$, then it follows from Eqs. (A.2) and (A.3) that h_i must be zero. Otherwise, the expected returns in the two equations cannot be equal.²⁸ The orthogonality between R_M and

²⁸ I assume here that $E[\text{HML}] \neq 0$. It is straightforward to show that the conditional expectation of HML cannot be zero, and there is no reason that the unconditional expectation should be zero.

HML establishes that $b_i = b_i'$. Also, M and HML span the tangency portfolio, so a_i'' is zero in the conditional regression (e.g., Shanken, 1987)

$$R_i(t) = a_i'' + b_{i,t}'' R_M(t) + h_{i,t}'' \text{HML}(t) + e_i''(t), \quad (\text{A.4})$$

where the conditional market beta and loading on HML are given by $b_{i,t}''$ and $h_{i,t}''$, respectively.

Because changes in the parameters are uncorrelated with the factor expected returns, $a_i = a_i' =$

0. It follows that h_i must be zero.

This proof depends on the assumption that time-variation in the conditional factor loadings, $b_{i,t}''$ and $h_{i,t}''$, is uncorrelated with the factors' expected returns. Although that assumption will not hold in general, it seems reasonable in my context because I am interested in the loadings changing over time with firm-specific variables, like B/M, not with macroeconomic conditions. Even if the assumption is not strictly true, previous studies suggest that correlation between the loadings and the factors has little effect on unconditional tests (e.g., Shanken, 1990). To gain some additional intuition, it may be useful to provide an example in which the assumption – and proposition – are exactly true. The example is similar to one presented by MacKinlay (1995).

The simplest scenario in which the assumption holds is when the factor expected returns are constant. To construct this example, assume first that the expected excess return on the market is constant (or, more generally, is uncorrelated with deviations from the CAPM). This assumption can be satisfied easily, regardless of whether the CAPM holds or not. We want to show that the expected return on HML can also be constant. Suppose that the residual covariance matrix, Σ , in market-model regressions is constant over time, non-singular, and has identical variances and covariances for all assets. For example, Sharpe's (1963) diagonal model or Ross's (1976) strict factor structure, with the market return as the only factor, are special cases if there are common residual variances. Also, suppose that conditional deviations from the CAPM, given by the $N \times 1$ vector α_t , have cross-sectional mean zero and variance σ_α^2 at

every t . The deviation for a given asset can fluctuate randomly over time, with mean zero, but we require that the cross-sectional dispersion is constant. These assumptions are stronger than necessary, but deliver the desired result, as we now show.

From MacKinlay (1995), the portfolio weights of HML are given by $w_{\text{HML},t} = \Sigma^{-1} \alpha_t$.²⁹ Notice that an asset's weight in the portfolio fluctuates over time with its mispricing. In this example, and in the Fama and French (1993) three-factor model, HML is a zero-investment portfolio, $\mathbf{1}' w_{\text{HML},t} = 0$. This fact can be derived from the requirement that $\mathbf{1}' \alpha_t = 0$ and the assumptions about Σ . Since HML is uncorrelated with the market return, its expected return equals $\alpha' w_{\text{HML},t} = \alpha_t' \Sigma^{-1} \alpha_t = \sum_i \sum_j \sigma_{ij}^{-1} \alpha_{i,t} \alpha_{j,t}$, where σ_{ij}^{-1} are the elements of the matrix Σ^{-1} . By assumption, Σ^{-1} has constant diagonals ($\sigma_{ii}^{-1} = c_1$) and constant off-diagonals ($\sigma_{ij}^{-1} = c_2$, for $i \neq j$). Therefore,

$$\begin{aligned} E_t[\text{HML}] &= c_1 \sum_i \alpha_{i,t}^2 + c_2 \sum_i \sum_{j \neq i} \alpha_{i,t} \alpha_{j,t} \\ &= N c_1 \sigma_\alpha^2 + c_2 \sum_i \alpha_{i,t} \sum_{j \neq i} \alpha_{j,t} \\ &= N c_1 \sigma_\alpha^2 + c_2 \sum_i \alpha_{i,t} (-\alpha_{i,t}) \\ &= N (c_1 - c_2) \sigma_\alpha^2, \end{aligned}$$

where the final three lines all use the fact that $\alpha_{i,t}$ has cross-sectional mean zero. In this example, the expected return on HML depends only on the total amount of mispricing, measured by the cross-sectional dispersion of $\alpha_{i,t}$. Therefore, the expected return on HML is constant, and more importantly, the unconditional factor loading, h , must be zero.

The example helps develop some intuition about the proposition. In general, it seems reasonable to believe that the expected return on HML will depend primarily on the cross-sectional dispersion in $\alpha_{i,t}$. When mispricing is large, investors can do much better than the

²⁹ To be precise, the vector of weights defined here does not guarantee that HML is orthogonal to R_M . Without loss of generality, I will assume that they are uncorrelated, but that is for convenience only.

CAPM, and HML has a high expected return. There does not seem to be any obvious reason that a given asset's mispricing should be strongly correlated with the overall amount mispricing in the market. If so, then the proposition should hold fairly well.

A.2. Factors

The factors used in this study are similar to those of Fama and French (1993), with a few minor differences. The three-factor model consists of market, size, and book-to-market factors. The market factor equals the return on the CRSP value-weighted index minus the T-bill rate at the beginning of the month. This factor differs somewhat from the market factor used by Fama and French, since they used only stocks with Compustat data to calculate the market return. However, there is little reason to limit the regression to stocks on Compustat, so all firms on CRSP are used for both the dependent portfolios and the market factor.

The size and book-to-market factors are calculated as follows. Each month, all stocks with market value data on CRSP for the previous month and book value data on Compustat for the previous fiscal year are sorted independently on size and B/M. I do not assume that book data become known until five months after fiscal year end. Following Fama and French, I define book equity as the book value of stockholder's equity minus the book value of preferred stock plus balance-sheet deferred taxes and investment tax credits, where the book value of preferred stock is given by redemption, liquidation, or par value, in that order of availability. Only firms with non-negative book equity and stock classified as common equity by CRSP are included.

Stocks are sorted into two size portfolios and three book-to-market portfolios, using as breakpoints the median market value and the 30th and 70th book-to-market percentiles of NYSE stocks, respectively. I calculate value-weighted returns for each of the six portfolios formed by the intersection of the two size and three book-to-market portfolios. In other words, returns are calculated for three portfolios of small stocks, with low, medium, and high B/M ratios, and for three portfolios of 'big' stocks, also with low, medium, and high B/M ratios. The

size factor, SMB, equals the average return on the three small portfolios minus the average return on the three big portfolios. The book-to-market factor, HML, equals the average return on the two high-B/M portfolios minus the average return on the two low-B/M portfolios. Hence, SMB and HML are returns on zero-investment portfolios designed to capture risk factors related to size and B/M, respectively.

A.3. Bootstrap simulations

The OLS slope estimate is biased upward in a regression of stock returns on lagged B/M (see Stambaugh, 1986). Since the bias in SUR estimates is unknown, I rely on bootstrap simulations to assess their sampling distribution.

The return regression can be thought of as part of the system

$$R_i(t) = \gamma_{i0} + \gamma_{i1} B/M_i(t-1) + e_i(t), \quad (\text{A.5})$$

$$B/M_i(t) = c_i + p_i B/M_i(t-1) + u_i(t). \quad (\text{A.6})$$

The bias in the OLS estimate of γ_{i1} is a function of p_i and $\text{cov}(e_i, u_i)$. Therefore, to estimate the bias in the SUR estimates, the simulation maintains the strong autocorrelation in B/M and the negative covariance between e_i and u_i that are observed in the data. Also, since SUR jointly estimates the system of equations for all portfolios, the simulation incorporates cross-sectional correlation among the residuals.

The bootstrap generates artificial time series of excess returns and B/M from eqs. (A.5) and (A.6). To construct returns, γ_{i0} is set equal to portfolio i 's average return and γ_{i1} is set equal to zero. Notice that the OLS bias is not a function of γ_{i1} (see eq. 2.8 in the text), so the value of γ_{i1} that is chosen should not be important. To construct B/M, the beginning value is given by the historical starting value and c_i and p_i are set equal to the sample estimates. The artificial time series, for 368 months, are then generated by sampling from the OLS residuals of the system, obtained after adjusting for the OLS bias in γ_{i1} . Each month of the sample, OLS produces a

vector of residuals from both equations, where the vectors are made up of the error terms for all portfolios. I randomly select, with replacement, pairs of residual vectors from this population.

Given these series, I estimate the return equations using the SUR methodology. The process is repeated 1500 times to construct an empirical distribution of SUR estimates. Since γ_{i1} equals zero by construction, the mean of the distribution estimates the bias in the SUR estimates. The covariance matrix provides an estimate of the SUR standard errors and covariances.

On the Predictability of Stock Returns: Theory and Evidence

Appendix B

This appendix supplements Chapter 3. We describe the Bayesian inference problem for the numerical simulations presented in Table 3.1.

The simulations are based on the renewal model of estimation risk, in which the mean of the dividend process is subject to periodic shocks. Dividends are assumed to follow a geometric random walk with a time-varying growth rate:

$$\ln d_{t+1} = g_k + \ln d_t + \varepsilon_{t+1}, \quad (\text{B.1})$$

where $\varepsilon_{t+1} \sim N[0, \sigma^2]$ and g_k is randomly drawn every K periods from a normal distribution with mean g^* and variance σ^2/s . At the beginning of a regime, investors' prior beliefs about g_k are $N[g^*, \sigma^2/h]$. After t periods in a regime ($t \leq K$), investors beliefs about g_k are $N[c_t, \sigma_{c,t}^2]$, where

$$c_t = \frac{h}{t+h} g^* + \frac{t}{t+h} \frac{1}{t} \sum_{i=1}^t \Delta \ln d_i, \quad (\text{B.2})$$

$$\sigma_{c,t}^2 = \frac{1}{t+h} \sigma^2. \quad (\text{B.3})$$

The predictive belief about log dividends next period is normally distributed with mean $c_t + \ln d_t$ and variance $[(t+h+1)/(t+h)]\sigma^2$. Actual dividends are log-normally distributed. Converting the expectations about log dividends into actual dividends, and extending the results to any dividend in the next q periods, where $t + q \leq K$ (that is, dividends in the current regime), we have that the predictive distribution of dividends is log-normal with mean

$$E_t^s [d_{t+q}] = d_t \exp \left[c_t + \frac{1}{2} q \sigma^2 + \frac{1}{2} q^2 \sigma_{c,t}^2 \right]. \quad (\text{B.4})$$

This equation fully takes into account the fact that changes in log dividends are correlated with changes in beliefs about the growth rate. In other words, investors recognize that their beliefs, both the mean and the variance, will evolve over time. After the end of the current regime,

investors expect dividends to grow once again at the rate g^* , and the variance of the growth rate is σ^2/h . Therefore, to derive beliefs about long-run dividends requires two steps: first, take the expectation conditional on the realized dividend at the end of the current regime, d_k , and then take the expectation conditional only on the current dividend, d_t . Details available on request.