

NBER WORKING PAPER SERIES

ESTIMATION RISK, MARKET EFFICIENCY,
AND THE PREDICTABILITY OF RETURNS

Jonathan Lewellen
Jay Shanken

Working Paper 7699
<http://www.nber.org/papers/w7699>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2000

We are grateful to Greg Bauer, John Campbell, John Cochrane, Kent Daniel, Gur Huberman, John Long, Bill Schwert, Robert Stambaugh, René Stulz, Jiang Wang, an anonymous referee, and seminar participants at Columbia University, Emory University, National Bureau of Economic Research, University of Michigan, University of Rochester, Washington University, Yale School of Management, and the American, European, and Northern Finance Association meetings for their helpful comments and suggestions. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 2000 by Jonathan Lewellen and Jay Shanken. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Estimation Risk, Market Efficiency, and the Predictability of Returns
Jonathan Lewellen and Jay Shanken
NBER Working Paper No. 7699
May 2000
JEL No. C11, D83, G12, G14

ABSTRACT

In asset pricing, estimation risk refers to investor uncertainty about the parameters of the return or cashflow process. We show that with estimation risk the observable properties of prices and returns can differ significantly from the properties perceived by rational investors. In particular, parameter uncertainty will tend to induce return predictability in ways that resemble irrational mispricing, and prices can violate familiar volatility bounds when investors are rational. Cross-sectionally, expected returns deviate from the CAPM even if investors attempt to hold mean-variance efficient portfolios, and these deviations can be predictable based on past dividends and prices. In short, estimation risk can be important for characterizing and testing market efficiency.

Jonathan Lewellen
Sloan School of Management
MIT
50 Memorial Drive, E52-436
Cambridge, MA 02142
(617) 258-8408
lewellen@mit.edu

Jay Shanken
William E. Simon Graduate School of
Business Administration
University of Rochester
Rochester, NY 14627
and NBER
(716) 275-4896
shanken@ssb.rochester.edu

Estimation risk, market efficiency, and the predictability of returns

1. Introduction

There is now 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) find that interest rates, dividend yield, aggregate book-to-market, and the default premium forecast 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) show that past returns also contain information about expected returns. In sum, there seems little doubt that expected returns vary both cross-sectionally and over time.¹

The interpretation of predictability, however, is more contentious. The empirical results 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 of asset-pricing questions.

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

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.

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 paper, we argue that there is a third potential source of return predictability: *estimation risk*. In the asset-pricing literature, estimation risk refers to investor uncertainty about the 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 consequence, estimation risk affects equilibrium prices and expected returns, and further, we show that it can be a source of predictability in a way that differs from other models with rational investors.

The theoretical literature on estimation risk 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 cash flows, and is therefore relevant for investment decisions.² 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

² 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

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. 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 cash flows. 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 cash flows diverge from the true values, asset prices deviate from their values in the absence of estimation risk.

estimate expected returns.

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 6). 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 cash flows. 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 the implications for market efficiency and the cross-section of expected returns.

The paper is organized as follows. Sections 2 and 3 introduce the basic model and derive capital market equilibrium. Section 4 examines the time-series properties of prices and returns and Section 5 explores the cross-sectional behavior of returns. Section 6 generalizes the model to incorporate informative priors, time-varying parameters, and non-stationary dividends, and presents simulation evidence from the general model. Section 7 concludes.

2. 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.

2.1. 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.

2.2. 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):

$$d_t \sim \text{MVN}[\delta, \Sigma], \quad (1)$$

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 6 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.

2.3. 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

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

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

$$x_t^* = \frac{1}{2\gamma} \left[\text{var}_t^s(p_{t+1} + d_{t+1}) \right]^{-1} \left[E_t^s(p_{t+1} + d_{t+1}) - (1+r)p_t \right], \quad (3)$$

where E_t^s and var_t^s denote the subjective expectation and variance at t .³ 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). 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

$$p_t = \frac{1}{1+r} \left[E_t^s(p_{t+1} + d_{t+1}) - 2\gamma \text{var}_t^s(p_{t+1} + d_{t+1}) \mathbf{1} \right]. \quad (4)$$

³ Throughout the paper we denote subjective moments with an 's' superscript.

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

3. 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 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.

3.1. 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. (4) and solving for price yields

$$p = \frac{1}{r} \delta - \frac{2\gamma}{r} \Sigma \mathbf{1}. \quad (5)$$

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

$$p_M = \iota' p = \frac{1}{r} \delta_M - \frac{2\gamma}{r} \sigma_M^2, \quad (6)$$

where the dividend on the market portfolio has expectation $\delta_M = \iota' \delta$ and variance $\sigma_M^2 = \iota' \Sigma \iota$. 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. (6) when $N = 1$.

3.2. *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

$$d_{t+1} \sim^s MVN\left[\bar{d}_t, \frac{t+1}{t} \Sigma\right]. \quad (7)$$

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. (7), it is clear that the subjective distribution of dividends, and consequently of 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. (4), yielding⁴

$$p_t = \frac{1}{r} \bar{d}_t - 2 \gamma \left[\sum_{k=1}^{\infty} \frac{1}{(1+r)^k} E_t^s \text{var}_{t+k-1}^s (p_{t+k} + d_{t+k}) \right] \cdot \iota. \quad (8)$$

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. (8). Variation in prices is driven entirely by the first term. Therefore, the subjective variance of returns is

$$\text{var}_t^s (p_{t+1} + d_{t+1}) = \left[1 + \frac{1}{r(t+1)} \right]^2 \left(\frac{t+1}{t} \right) \Sigma. \quad (9)$$

Substituting into eq. (8) yields the equilibrium pricing function:

$$p_t = \frac{1}{r} \bar{d}_t - 2 \gamma f(t) \Sigma \iota, \quad (10)$$

where

$$f(t) = \sum_{k=1}^{\infty} \frac{1}{(1+r)^k} \left[1 + \frac{1}{r(t+k)} \right]^2 \left(\frac{t+k}{t+k-1} \right). \quad (11)$$

The equilibrium price is similar to the price with perfect information (eq. 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

⁴ Eq. (8) imposes the transversality condition $\lim_{k \rightarrow \infty} E_t[p_{t+k}]/(1+r)^k = 0$, which will be satisfied in equilibrium.

term. The function $f(t)$ decreases as t gets larger and converges to $1/r$ in the 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 6 we allow the true parameters to change, so that investors never completely learn the dividend process.

We noted in Section 2 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

$$p_{M,t} = \iota' p_t = \frac{1}{r} \bar{d}_{M,t} - 2 \gamma f(t) \sigma_M^2, \quad (12)$$

where $\bar{d}_{M,t} = \iota' \bar{d}_t$ 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. (12) when $N = 1$.

Several colleagues have noted that the pricing function in eq. (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.

4. 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. (6) and (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

$$p_{t+1} - p_t = \frac{1}{r}(\bar{d}_{t+1} - \bar{d}_t) + 2\gamma[f(t) - f(t+1)]. \quad (13)$$

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.

4.1. 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 cash flows 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

$$R_{t+1} = d_{t+1} + \frac{1}{r(t+1)}(d_{t+1} - \bar{d}_t). \quad (14)$$

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:

$$E_t^s[p_{t+1} - p_t] = 0. \quad (15)$$

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. (14), the true conditional expected return is

$$E_t[R_{t+1}] = \delta + \frac{1}{r(t+1)}(\delta - \bar{d}_t). \quad (15)$$

It is clear that R_{t+1} is negatively related to past dividends.⁵ Although dividends are IID by assumption, price revisions are negatively correlated with past cash flows. 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 cash flows.

From eq. (15), prices, dividends, and returns all predict time-variation in expected returns. For

⁵ For simplicity, we examine the predictability of *gross* returns rather than *rates of* return. The analysis with rates of return is more difficult because it involves expectations of ratios, but the qualitative results are similar.

example, suppose we are interested in the autocovariance of returns:⁶

$$\text{cov}[R_t, R_{t+1}] = -\frac{1}{r t(t+1)} \sigma^2. \quad (16)$$

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,

$$\text{cov}[d_t, R_{t+1}] = -\frac{1}{r t(t+1)} \sigma^2. \quad (17)$$

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. 1 depicts a sample price path for the risky asset. The figure assumes that investors are risk-neutral and the riskless rate is 0.05. Dividends have mean 0.05 and standard deviation 0.10, taken to be similar to the dividend yield and 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%.⁷

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

⁶ This covariance is time-dependent because estimation risk declines over time. We will break the strong connection between time and predictability in Section 6 when we allow the true dividend process to change.

⁷ 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 6.

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. Similar results for actual stock market data would be interpreted as 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

$$E_t[R_{t+H}^H] = H\delta + \frac{H}{r(t+H)}(\delta - \bar{d}_t). \quad (18)$$

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

$$\text{cov}[R_t^H, R_{t+H}^H] = -\frac{H^2}{r t (t+H)} \sigma^2, \quad (19)$$

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 cash flows.

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. (14),

$$UR_{t+1} = \left[1 + \frac{1}{r(t+1)} \right] (d_{t+1} - \bar{d}_t). \quad (20)$$

It follows that

$$E_t[UR_{t+1}] = \left[1 + \frac{1}{r(t+1)} \right] (\delta - \bar{d}_t). \quad (21)$$

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. 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 6, 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 cash flows.

4.2. 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

$$\text{var}_t [p_{t+1}] = \left[\frac{1}{r(t+1)} \right]^2 \sigma^2, \quad (22)$$

and the unconditional variance is

$$\text{var}[p_{t+1}] = \frac{1}{r^2(t+1)} \sigma^2. \quad (23)$$

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 is 0.05, and dividends are distributed with mean 0.05 and standard deviation 0.10. (These are the values used in Fig. 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_t 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

$$\text{var}(p_t) \leq \text{var}(p_t^*), \quad (24)$$

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

$$p_t^* = \sum_{k=1}^{\infty} \frac{1}{(1+r)^k} d_{t+k} - 2\gamma f(t)\sigma^2. \quad (25)$$

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

$$\text{var}(p_t^*) = \sum_{k=1}^{\infty} \frac{1}{(1+r)^{2k}} \sigma^2 = \frac{1}{r^2 + 2r} \sigma^2. \quad (26)$$

Comparing this to eq. (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

$$\text{var}_t[R_{t+1}] = \left[1 + \frac{1}{r(t+1)} \right]^2 \sigma^2. \quad (27)$$

Comparing this to the variance of the subjective distribution (eq. 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.

5. 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.’⁸ To see why this is important, recall that unexpected returns in the model equal

$$UR_{t+1} = \left[1 + \frac{1}{r(t+1)} \right] (d_{t+1} - \bar{d}_t), \quad (28)$$

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 6.

5.1. 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

$$\text{var}_t [R_{t+1}] = \left[1 + \frac{1}{r(t+1)} \right]^2 \Sigma. \quad (29)$$

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

⁸ 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 $MVN[\delta^*, \Sigma/h]$, where h is a measure of prior information. The diffuse prior can be interpreted as the limiting distribution as h approaches zero.

covariance matrices. Comparing eqs. (9) and (29), we find that the subjective covariance matrix equals the true covariance matrix multiplied by $(t+1)/t$, and both are proportional to Σ .⁹

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

$$\beta = \frac{1}{\text{var}(R_{M,t})} \text{cov}(R_t, R_{M,t}) = \frac{1}{\mathbf{1}' \Sigma \mathbf{1}} \Sigma \mathbf{1}. \quad (30)$$

Note also that eq. (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.¹⁰

5.2. Expected returns and the CAPM

In Section 4, 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

$$E_t^s[R_{t+1}] = r p_t + \beta [E_t^s(R_{M,t+1}) - r p_{M,t}]. \quad (31)$$

Eq. (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

⁹ 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 6.

¹⁰ 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

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 the CAPM, given by

$$a_{t+1} = R_{t+1} - r p_t - \beta [R_{M,t+1} - r p_{M,t}]. \quad (32)$$

Note that a_{t+1} is similar to the vector of unexpected returns, except that the realized return on the market enters eq. (32) rather than the expected market return. We know from Section 4 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:

$$E_t^s [a_{t+1}] = 0. \quad (33)$$

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. (32) and taking expectations yields

$$E_t [a_{t+1}] = - \left[1 + \frac{1}{r(t+1)} \right] [(\bar{d}_t - \delta) - (\bar{d}_{M,t} - \delta_M) \beta]. \quad (34)$$

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

by $p_{m,t} / p_{i,t}$. Rate-of-return betas will change unless relative prices remain the same, which will not be true in general.

any asset i :

$$\begin{aligned}\text{cov}[p_{i,t}, a_{i,t+1}] &= -\left[1 + \frac{1}{r(t+1)}\right] \frac{1}{r t} [\text{var}(d_i) - \beta_i \text{cov}(d_i, d_M)] \\ &= -\left[1 + \frac{1}{r(t+1)}\right] \frac{1}{r t} \text{var}(\varepsilon_i),\end{aligned}\quad (35)$$

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. (35) is essentially a time-series relation. The predictability of a_{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

$$\text{cov}_{t+1}^{\text{cs}}[p_{i,t}, a_{i,t+1}] = \frac{1}{N} \sum_i (a_{i,t+1} - \bar{a}_{t+1}^{\text{cs}}) (p_{i,t} - \bar{p}_t^{\text{cs}}). \quad (36)$$

Taking the unconditional expectation yields

$$E[\text{cov}_{t+1}^{\text{cs}}(p_{i,t}, a_{i,t+1})] = \frac{1}{N} \sum_i \text{cov}(a_{i,t+1}, p_{i,t}) < 0, \quad (37)$$

which is negative because every covariance term is negative (see eq. 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. (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 4, 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, expected 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.

6. Informative priors, time-varying parameters, and simulation evidence

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.

6.1. 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

$$d_{t+1} \sim^s N \left[\frac{h}{t+h} \delta^* + \frac{t}{t+h} \bar{d}_t, \frac{t+h+1}{t+h} \sigma^2 \right]. \quad (38)$$

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 prior beliefs as well as the information in realized dividends. Denote the mean of the subjective distribution as m_t . At time t , the price of the risky asset equals

$$\begin{aligned} p_t &= \frac{1}{r} m_t - 2 \gamma f(t+h) \sigma^2 \\ &= \frac{1}{r} \frac{t}{t+h} \bar{d}_t + \frac{1}{r} \frac{h}{t+h} \delta^* - 2 \gamma f(t+h) \sigma^2, \end{aligned} \quad (39)$$

where $f(t)$ is defined in eq. (11). With informative priors, the price contains a new term corresponding to the initial belief about expected dividends. It is clear from eq. (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,

$$\text{cov}[p_t, R_{t+h}] = -\frac{t}{r^2(t+h)^2(t+h+1)}\sigma^2, \quad (40)$$

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 6.3, 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

$$m_t = [\Omega^{-1} + t \Sigma^{-1}]^{-1} [\Omega^{-1} \delta^* + t \Sigma^{-1} \bar{d}_t], \quad (41)$$

$$\Pi_t = [\Omega^{-1} + t \Sigma^{-1}]^{-1}. \quad (42)$$

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 predictive 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

$$m_t = \frac{h}{t+h} \delta^* + \frac{t}{t+h} \bar{d}_t, \quad (43)$$

$$\Pi_t = \frac{1}{t+h} \Sigma. \quad (44)$$

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

$$p_t = \frac{1}{r} m_t - 2 \gamma v(t) \mathbf{1}, \quad (45)$$

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

$$E[\text{cov}_{t+1}^{\text{cs}}(p_{i,t}, a_{i,t+1})] = \left[1 + \frac{1}{t+h+1}\right] \frac{t}{r(t+h)^2} [h \text{var}^{\text{cs}}(\delta) - \overline{\text{var}}(\varepsilon_i)], \quad (46)$$

where $\text{var}(\varepsilon_i)$ is the residual variance when the asset's dividend is regressed on the market dividend (see eq. 35), and $\overline{\text{var}}(\varepsilon_i)$ 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 5 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 5, 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

$$\text{var}_t(\mathbf{R}_{t+1}) = \text{var}_t\left[\mathbf{d}_{t+1} + \frac{1}{r} \mathbf{m}_{t+1}\right] = \text{var}[\mathbf{M}_{t+1} \mathbf{d}_{t+1}] = \mathbf{M}_{t+1} \boldsymbol{\Sigma} \mathbf{M}'_{t+1}, \quad (47)$$

where

$$\mathbf{M}_{t+1} = \mathbf{I} + \frac{1}{r} \boldsymbol{\Pi}_{t+1} \boldsymbol{\Sigma}^{-1}. \quad (48)$$

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:

$$\text{var}_t^s(\mathbf{R}_{t+1}) = \mathbf{M}_{t+1} [\boldsymbol{\Pi}_t + \boldsymbol{\Sigma}] \mathbf{M}_{t+1}^T. \quad (49)$$

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

6.2. 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 scenario 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.¹¹ 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

¹¹ For example, suppose the dividend mean δ_t follows a random walk, dividends have conditional variance σ^2 , and the shocks to δ_t are uncorrelated with dividends and have variance σ_s^2 . In the long-run, investors' beliefs about δ_t will be $N[m, \sigma_h^2]$, where σ_h^2 is time-invariant and $\sigma_h^2 > \sigma_s^2$.

short-run mean, δ_k . As discussed above, investors' priors at the beginning 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

$$m_t = \frac{h}{t+h} \delta^* + \frac{t}{t+h} \bar{d}_t, \quad (50)$$

identical to eq. (38). Thus, price equals

$$p_t = AF_{K-t} m_t + \frac{1}{(1+r)^{K-t}} \frac{\delta^*}{r}, \quad (51)$$

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

$$UR_{t+1} = \left[1 + \frac{AF_{K-t-1}}{t+h+1} \right] (d_{t+1} - m_t). \quad (52)$$

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

$$E_t[UR_{t+1}] = \left[1 + \frac{AF_{K-t-1}}{t+h+1} \right] (\delta_k - m_t), \quad (53)$$

which follows immediately from eq. (52). As in our basic model, the true unexpected return is negatively

related to past dividends and prices. Consequently, *taking the value of δ_k as given*, the covariance between excess returns and lagged prices equals

$$\begin{aligned}\text{cov}[p_t, \text{UR}_{t+1}] &= -\text{AF}_{K-t} \left(1 + \frac{\text{AF}_{K-t-1}}{t+h+1} \right) \text{var}(m_t) \\ &= -\text{AF}_{K-t} \left(1 + \frac{\text{AF}_{K-t-1}}{t+h+1} \right) \left(\frac{t}{t+h} \right)^2 \frac{1}{t} \sigma^2,\end{aligned}\tag{54}$$

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. (54).¹² Specifically, the unconditional covariance equals

$$\begin{aligned}\text{cov}[p_t, \text{UR}_{t+1}] &= \text{AF}_{K-t} \left(1 + \frac{\text{AF}_{K-t-1}}{t+h+1} \right) [\text{cov}(d_{t+1}, m_t) - \text{var}(m_t)], \\ &= \text{AF}_{K-t} \left(1 + \frac{\text{AF}_{K-t-1}}{t+h+1} \right) \frac{t}{(t+h)s} \left[1 - \frac{t+s}{t+h} \right] \sigma^2.\end{aligned}\tag{55}$$

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 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,

¹² In statistical terms, the expected conditional covariance does not equal the unconditional covariance because the conditional means of the variables move together over time.

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.¹³ 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 the 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.¹⁴ 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 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. However, it is not common to include regime dummies in predictive regressions, nor is it easy to identify regime changes. Rather than speculate further, we use simulations to explore the sample covariance with renewal.

6.3. Simulations

¹³ 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.

¹⁴ 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 δ_k* because the true correlation is negative.

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

$$\ln d_{t+1} = g_k + \ln d_t + \varepsilon_{t+1}, \quad (56)$$

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. The appendix describes the Bayesian inference problem for this model.

Table 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.¹⁵ An important complication arises because the slope coefficient in these regressions suffers from a significant small-sample bias (see Stambaugh, 1999, for a thorough discussion). 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).¹⁶ In addition, the table

¹⁵ We focus on predictability here, but other moments of the return distribution are also affected by estimation risk. This might be a useful area of future research.

¹⁶ To derive the bias, Stambaugh 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

reports results when investors perfectly observe the dividend process. The difference between the bias-adjusted coefficients with and without perfect information gives an estimate of the predictability caused by estimation risk.

Table 1 shows that estimation risk can induce predictability even in steady state. The results suggest that the negative conditional covariance between returns and lagged prices 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, are of particular interest because they show roughly how prices behave before we reach steady state (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 at least 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

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 1.

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 realize. 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.

7. 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 dividends and prices, and can actually be negative at times. 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. This is probably not an easy task: it would necessarily require some judgment about what constitutes a ‘reasonable’ prior and an examination of rational forecasts generated by a range of such priors. Fama (1971) has emphasized that empirical tests of market efficiency actually test a joint hypothesis of market efficiency and an assumed model of expected returns. Our study suggests that empirical tests also require an additional assumption about investors’ prior beliefs. Although the role of prior beliefs and learning is typically ignored, these considerations may be just as important in disentangling behavioral and rational explanations for empirical anomalies.

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.

Appendix

This appendix describes the Bayesian inference problem in the renewal model. Dividends are assumed to follow a geometric random walk with a time-varying growth rate:

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

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

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

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

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

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

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.

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.
- Barberis, Nicholas, 1999, Investing for the long run when returns are predictable, forthcoming in *Journal of Finance*.
- 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).
- Black, Fischer, 1986, Noise, *Journal of Finance* 41, 529-543.
- Bossaerts, Peter, 1997, The dynamics of equity prices in fallible markets, Working paper (California Institute of Technology, Pasadena, California).
- Brennan, Michael and Y. 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.
- 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.
- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and over-reactions, *Journal of Finance* 53, 1839-1885.
- 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, 1970, Efficient capital markets: A review of theory and empirical work, *Journal of Finance* 25, 383-417.
- Fama, Eugene, 1976, *Foundations of Finance* (Basic Books, New York, NY).
- Fama, Eugene, 1991, Efficient capital markets: II, *Journal of Finance* 46, 1575-1617.
- Fama, Eugene, 1998, Market efficiency, long-term returns, and behavioral finance, *Journal of Financial Economics* 49, 283-306.
- 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 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.
- Harvey, Campbell, 1989, Time-varying conditional covariances in tests of asset pricing models, *Journal of Financial Economics* 24, 289-317.
- 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.
- Jorion, Philippe, 1985, International portfolio decisions with estimation risk, *Journal of Business* 58, 259-278.
- 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.
- 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.
- 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.
- 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.
- Poterba, James and Lawrence Summers, 1988, Mean reversion in stock prices: Evidence and implications, *Journal of Financial Economics* 22, 27-59.
- Shanken, Jay, 1990, Intertemporal asset pricing: An empirical investigation, *Journal of Econometrics* 45, 99-120.
- Shiller, Robert, 1981, Do stock prices move too much to be justified by subsequent changes in dividends?, *American Economic Review* 7, 421-436.
- Stambaugh, Robert, 1997, Analyzing investments whose histories differ in length, *Journal of Financial Economics* 45, 285-331.
- Stambaugh, Robert, 1999, Predictive regressions, Working paper (Wharton School, University of Pennsylvania, Philadelphia, PA).
- 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.
- Williams, Joseph, 1977, Capital asset prices with heterogeneous beliefs, *Journal of Financial Economics* 5, 219-239.
- Zellner, Arnold, 1971, *An Introduction to Bayesian Inference in Econometrics* (John Wiley and Sons, New York, NY).

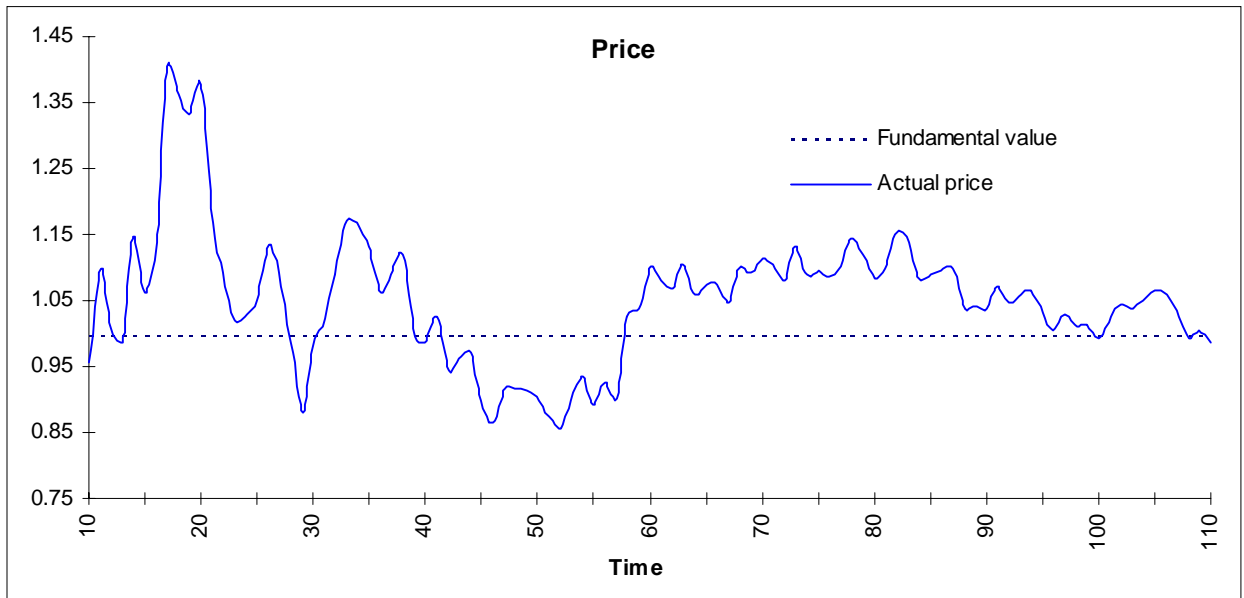


Figure 1
Equilibrium price of the risky asset

This figure illustrates a sample price path for the risky asset when the dividend process is known ('fundamental value' in the figure; see eq. 6 in the text) and when investors must estimate expected dividends ('actual price'; see eq. 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.

Table 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 results for 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 (approximate because 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).

		s	Estimation risk			Perfect information			Difference		
			h			h			h		
			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	0.78	0.32	-0.53	-0.28	-0.29	-0.16	1.06	0.60	-0.37
		25	1.16	0.84	0.15	-0.28	-0.34	-0.11	1.44	1.18	0.25
		49	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	0.29	-0.56	-2.95	-0.27	-0.33	-0.29	0.56	-0.23	-2.66
		25	1.06	0.26	-1.36	-0.29	-0.31	-0.34	1.35	0.57	-1.01
		49	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	0.16	-1.35	-4.15	-0.33	-0.26	-0.34	0.49	-1.08	-3.81
		25	1.22	0.14	-1.95	-0.27	-0.31	-0.34	1.49	0.45	-1.61
		49	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