

Estimating the Equity Premium

R. Glen Donaldson, Mark J. Kamstra, and Lisa A. Kramer*

Abstract

Existing empirical research investigating the size of the equity premium has largely consisted of a series of innovations around a common theme: producing a better estimate of the equity premium by using better data or a better estimation technique. The equity premium estimate that emerges from most of this work matches one moment of the data alone: the mean difference between an estimate of the return to holding equity and a risk-free rate. We instead match multiple moments of U.S. market data, exploiting the joint distribution of the dividend yield, return volatility, and realized excess returns, and find that the equity premium lies within 50 basis points of 3.5%, a range much narrower than was achieved in previous studies. Additionally, statistical tests based on the joint distribution of these moments reveal that only those models of the conditional equity premium that embed time variation, breaks, and/or trends are supported by the data. In order to develop the joint distribution of the dividend yield, return volatility, and excess returns, we need a model of price and return fundamentals. We document that even recently developed analytically tractable models that permit autocorrelated dividend growth rates and discount rates impose restrictions that are rejected by the data. We therefore turn to a wider range of models, requiring numerical solution methods and parameter estimation by the simulated method of moments.

*Donaldson, glen.donaldson@sauder.ubc.ca, Sauder School of Business, University of British Columbia, 2053 Main Mall, Vancouver, BC V6T 1Z2, Canada; Kamstra, mkamstra@yorku.ca, Schulich School of Business, York University, 4700 Keele St., Toronto, ON M3J 1P3, Canada; and Kramer, lisa.kramer@rotman.utoronto.ca, Rotman School of Management, University of Toronto, 105 St. George St., Toronto, ON M5S 3E6, Canada. We are grateful for the suggestions of Stephen Brown (the editor), Wayne Ferson, Mark Fisher, Ian Garrett, Joel Hasbrouck, Robert Hodrick (the referee), Raymond Kan, Patrick Kelly, Alan Kraus, Tom McCurdy, Federico Nardari, Cesare Robotti, Jacob Sagi, and Tan Wang; participants of the Western Finance Association Meetings, the Northern Finance Association Meetings, the Canadian Econometrics Study Group, the European Econometric Society Meetings, and the investment conference of the University of Colorado at Boulder's Burrige Center for Securities Analysis and Valuation; and seminar participants at the Board of Governors of the Federal Reserve System, Emory University, the Federal Reserve Bank of Atlanta, Queen's University, the U.S. Securities and Exchange Commission, and the University of British Columbia. We thank the Social Sciences and Humanities Research Council of Canada for financial support. Any remaining errors are our own. Previous versions of this paper were titled "Stare Down the Barrel and Center the Crosshairs: Targeting the Ex Ante Equity Premium" and "Estimating the Ex Ante Equity Premium."

I. Introduction

Financial economic theory is often concerned with the premium that investors anticipate *ex ante*, when they first decide whether to purchase risky stocks instead of risk-free debt. In contrast, since it is difficult to directly observe what investors anticipate, empirical tests of the equity premium traditionally focus on the returns that investors received *ex post*.¹ It is well known that estimates of the equity premium based on *ex post* data can be very imprecise; such estimates have very wide margins of error, as wide as 1,000 basis points (bp) in typical studies and 320 bp in some recent studies. This imprecision makes it challenging to employ equity premium estimates for common practical purposes, including evaluating the equity premium puzzle, performing valuation, conducting capital budgeting, and even determining if the equity premium has changed over time.

In this paper we focus on the development of a more precise estimate of the equity premium. Our approach involves identifying a range of values and conditional models of the equity premium that are consistent with the *joint* distribution of several key financial statistics observed in U.S. market data, including dividend growth rates, interest rates, Sharpe ratios, price-dividend ratios, volatilities, and of course, excess returns. It is the consideration of *multiple* moments of financial data that delivers increased precision and new insights into the equity premium process.

Our results suggest that the equity premium lies within 50 bp of 3.5%. We also find that equity premium models that allow for time variation, breaks, and/or trends are the models that best match the experience of U.S. markets and are the only models not rejected by our specification tests. This suggests that time variation, breaks, and/or trends are critical features of the equity premium process.

The recent empirical literature that examines the equity premium has largely consisted of a series of innovations around a common theme: producing a better estimate of the unconditional equity premium by estimating either the mean equity return or the risk-free rate (or both), using novel (arguably better) data, or by using a better (more efficient) estimation technique. Our work fits squarely in this tradition, exploiting multiple moments of the data to provide a more precise estimate of the equity premium. Previous work in the area has included insights such as exploiting dividend yields or earnings yields to provide new, more precise estimates of the expected return to holding stocks (Fama and French (2002)); looking at stock returns across many countries to account for survivorship issues (Jorion and Goetzmann (1999)); looking across many countries to decompose the equity premium into dividend growth, price-dividend ratio, dividend yield, and real exchange rate components (Dimson, Marsh, and Staunton (2008)); modeling

¹The equity premium literature is large, continuously growing, and much too vast to fully cite here. For recent work, see Bansal and Yaron (2004), Graham and Harvey (2005), Barro (2005), and Bonaparte (2008). For excellent surveys, see Kocherlakota (1996), Siegel and Thaler (1997), Mehra and Prescott (2003), and Mehra (2003).

equity premium structural breaks in a Bayesian econometric framework (Pástor and Stambaugh (2001)); or computing out-of-sample forecasts of the distribution of excess returns, allowing for structural breaks that are identified in real time (Maheu and McCurdy (2009)). Most of this work estimates the equity premium by considering one moment of the data at a time, typically the mean difference between an estimate of the expected return to holding equity and a risk-free rate, though Maheu and McCurdy consider higher-order moments of the excess return distribution and Pástor and Stambaugh incorporate return volatility and direction of price movements by use of priors.

In spite of these advances, point estimates of the equity premium produced in the existing literature are still imprecise. There is also less of an emerging consensus than one would hope. Fama and French (2002) produce point estimates of $2.55\% \pm 1.60$ (2 standard deviations) using dividend yields and $4.32\% \pm 4.00$ using earnings yields. Pástor and Stambaugh (2001) estimate the equity premium at the end of the 1990s to be $4.8\% \pm 2.80$. Claus and Thomas (2001) estimate the equity premium to be *no more* than 3%. Welch (2000), surveying academic financial economists, estimates the consensus equity premium to be between 6% and 7% (depending on the horizon). Welch (2008) finds that academics have recently lowered their estimates to about 5%. Based on a survey of chief financial officers (CFOs), Graham and Harvey (2005) estimate the 10-year equity premium to be 3.66%. We believe that the lack of consensus across the literature is intimately tied to the imprecision of techniques typically used to estimate the equity premium, such as the simple average excess return. That is, the various estimates cited previously all fall within 2 standard errors of the sample mean estimate of the equity premium, based on U.S. data. Further, the studies that provide these estimates do not explicitly consider which models of the equity premium process can be rejected by observed data, although Pástor and Stambaugh's analysis strongly supports a model that incorporates breaks in the equity premium process.

Our approach necessitates the use of a model of price and return fundamentals. A traditional way forward is to employ a model with a closed-form solution. Some analytic approaches require constant dividend growth rates and discount rates (to calculate a fundamental stock price as the present value of expected future dividends as in Gordon (1962)), while others require log linearizations (such as the Campbell and Shiller (1988) solution method). Some recent developments allow autocorrelated covarying dividend growth rates and discount rates (exploiting affine models and joint normality of the logarithm of the discount rate and growth rate of dividends; see, e.g., Ang and Liu (2001), (2007)). We have explored the previous types of analytically tractable models and find without exception that these models are rejected by the data (producing results that are significantly at odds with various moments of the observed data). In light of this finding, we are unable to approach this exercise using methods that would provide closed-form solutions. We therefore employ a simulation-based numerical solution method. This allows us to incorporate, for instance, equity premiums with trends and breaks, discount rate processes that are not joint lognormally distributed with dividend growth rates, and virtually any other irregular feature of discount rates and dividend growth rates.

The numerical solution approach we adopt first builds on the fundamental valuation dividend-discounting method of Donaldson and Kamstra (DK) (1996). This technique permits the simulation of fundamental prices, returns, and return volatility for a given value of the equity premium.² Second, the simulated dividend yields, excess returns, etc., produced for a range of values of the assumed equity premiums, are matched to the realized U.S. data. Our estimate of the equity premium is that value that leads to the best match between the simulated moments and the moments of the observed data. This is the simulated method of moments (SMM)³ estimation technique. SMM forms estimates of model parameters by using a given model with a given set of parameter values to simulate moments of the data (e.g., means or volatilities), measuring the distance between the simulated moments and the realized data moments, and repeating with new parameter values until the parameter values that minimize the (weighted) distance are found.⁴ The parameter estimates that minimize this distance are consistent for the true values, are asymptotically normally distributed, and display the attractive feature of permitting tests that can reject misspecified models. The SMM technique has been described as “estimating on one group of moments, testing on another” (see Cochrane (2001), Sect. 11.6). We use SMM rather than generalized method of moments (GMM) because, as we show later, the economic model we employ includes latent variables and therefore cannot be solved without the use of SMM.

The remainder of our paper proceeds as follows. The basic methodology of our simulation approach to estimating equity premiums is presented in Section II, along with details on estimating the equity premium. In Section III we compare univariate financial statistics that arise in our simulations with U.S. market data, including dividend yields, return volatility, excess returns, and Sharpe ratios. Our results confirm that the simulations generate data broadly consistent with the U.S. market data and, taken one at a time, these financial statistics imply that the conditional equity premium has indeed been falling over time and that the unconditional equity premium lies in a range much narrower than between 2% and 8%. We determine how much narrower in Section IV by exploiting the full power of the simulation methodology. Section V concludes.

II. Methodology

Consider a stock for which the price P_t is set at the beginning of each period t and that pays a dividend D_{t+1} at the end of period t . The return to holding this

²DK (1996) find that if dividend growth rates and discount rates are allowed to be time-varying and dependent, as well as cross-correlated, the fundamental prices and returns that come out of dividend discounting match observed prices and returns, even during extreme events like stock market crashes.

³SMM was developed by McFadden (1989) and Pakes and Pollard (1989). Examples of papers that employ SMM in an asset pricing context are Duffie and Singleton (1993) and Corradi and Swanson (2005).

⁴We adopt the typical implementation of SMM to weight the moments inversely to their estimated precision; that is to minimize the product of the moments weighted by the inverse of the covariance matrix of the moments.

stock over t is defined as $R_{t+1} = (D_{t+1} + P_{t+1} - P_t)/P_t$. The risk-free rate, set at the beginning of each period, is denoted $r_{f,t}$. The unconditional equity premium, π , is defined as the difference between the unconditional expectation of the return on risky assets, $E[R_{t+1}]$, and the risk-free rate, $E[r_{f,t}]$:⁵

$$(1) \quad \pi \equiv E[R_{t+1}] - E[r_{f,t}].$$

We do not observe this equity premium. Empirically we observe the returns that investors receive ex post, after they have purchased the stock and held it over some period of time during which random economic shocks impact prices. Hence, the equity premium is typically estimated using historical equity returns and risk-free rates. Define \bar{R} as the average historical annual return on the Standard & Poor's (S&P) 500 and \bar{r}_f as the average historical return on U.S. T-bills. Then we can estimate the equity premium, $\hat{\pi}$, as

$$(2) \quad \hat{\pi} \equiv \bar{R} - \bar{r}_f.$$

There is no reason to believe that the stock return realized ex post is exactly the same as the return investors anticipated ex ante, hence even a 6% ex post equity premium in the U.S. data may not be a challenge to economic theory. Therefore we ask the following question: If investors' equity premium is π , what is the probability that the U.S. economy could randomly produce an ex post premium of at least 6%? The answer to this question has implications for whether or not the 6% ex post premium observed in the U.S. data is consistent with various values of π with which standard economic theory may be more compatible. We also ask a deeper question: If investors' equity premium is π , what is the probability that we would observe the various *combinations* of key financial statistics that have been realized in the U.S., such as high Sharpe ratios and low dividend yields, high return volatility and a high ex post equity premium, and so on? The analysis of multivariate distributions of these statistics allows us to narrow substantially the range of equity premiums consistent with the U.S. market data relative to previous studies that have considered only univariate distributions.

Because the joint distribution of the financial statistics we wish to consider is difficult or impossible to estimate accurately, in particular the joint distribution *conditional* on a given equity premium value, we use simulation techniques to estimate this distribution. The simulated joint distribution allows us to conduct formal statistical tests that a given equity premium could have produced the U.S. experience. Most of our models employ a conditionally varying equity premium, so that a simulation described as having an equity premium of 2.75% actually has an unconditional equity premium of 2.75%, while period by period the conditional equity premium can vary somewhat from this mean value.

⁵See, for instance, Mehra and Prescott ((1985), eq. (14)). We consider conditional equity premium models later.

A. Matching Moments

Consider the valuation of a stock. Define $1 + r_t$ as the gross rate investors use to discount payments received during period t . The fundamental price of the stock is then

$$(3) \quad P_t = E_t \left[\frac{D_{t+1} + P_{t+1}}{1 + r_t} \right],$$

where E_t is the conditional expectations operator incorporating information available to the market at the beginning of period t . Assuming the usual transversality condition, we can derive equation (4) by recursively substituting out for future prices in equation (3):

$$(4) \quad P_t = E_t \left[\sum_{j=0}^{\infty} \left(\prod_{i=0}^j \frac{1}{1 + r_{t+i}} \right) D_{t+j+1} \right].$$

Defining the growth rate of dividends over the period t as $g_{t+1} \equiv (D_{t+1} - D_t)/D_t$, we can rewrite equation (4) as follows:⁶

$$(5) \quad P_t = D_t E_t \left[\sum_{j=0}^{\infty} \left(\prod_{i=0}^j \left[\frac{1 + g_{t+i+1}}{1 + r_{t+i}} \right] \right) \right].$$

Hence we can rewrite equation (1) as

$$(6) \quad \pi \equiv E \left[\left(D_{t+1} + D_{t+1} E_{t+1} \left\{ \sum_{j=0}^{\infty} \prod_{i=0}^j \frac{1 + g_{t+i+2}}{1 + r_{t+i+1}} \right\} - D_t E_t \left\{ \sum_{j=0}^{\infty} \prod_{i=0}^j \frac{1 + g_{t+i+1}}{1 + r_{t+i}} \right\} \right) / \left(D_t E_t \left\{ \sum_{j=0}^{\infty} \prod_{i=0}^j \frac{1 + g_{t+i+1}}{1 + r_{t+i}} \right\} \right) - r_{f,t} \right]$$

⁶One attractive feature of expressing the present value stock price as in equation (5), in terms of dividend growth rates and discount rates, is that this form highlights the irrelevance of inflation, at least to the extent that expected and realized inflation are the same. Notice that working with nominal growth rates and discount rates, as we do, is equivalent to working with deflated nominal rates (i.e., real rates). That is, $[1 + (g_{t+1} - I_t)/(1 + I_t)] / [1 + (r_t - I_t)/(1 + I_t)] = (1 + g_{t+1})/(1 + r_t)$, where I_t is inflation. Working with nominal values in our simulations removes a potential source of measurement error associated with attempts to estimate inflation.

or

$$(7) \quad \pi \equiv E \left[\left((1 + g_{t+1}) \left(1 + E_{t+1} \left\{ \sum_{j=0}^{\infty} \Pi_{i=0}^j \frac{1 + g_{t+i+2}}{1 + r_{t+i+1}} \right\} \right) - E_t \left\{ \sum_{j=0}^{\infty} \Pi_{i=0}^j \frac{1 + g_{t+i+1}}{1 + r_{t+i}} \right\} \right) / \left(E_t \left\{ \sum_{j=0}^{\infty} \Pi_{i=0}^j \frac{1 + g_{t+i+1}}{1 + r_{t+i}} \right\} - r_{f,t} \right) \right].$$

In the case of a constant equity premium π , and a possibly time-varying risk-free interest rate, we can rewrite equation (7) as

$$(8) \quad \pi \equiv E \left[\left((1 + g_{t+1}) \left(1 + E_{t+1} \left\{ \sum_{j=0}^{\infty} \Pi_{i=0}^j \frac{1 + g_{t+i+2}}{1 + \pi + r_{f,t+i+1}} \right\} \right) - E_t \left\{ \sum_{j=0}^{\infty} \Pi_{i=0}^j \frac{1 + g_{t+i+1}}{1 + \pi + r_{f,t+i}} \right\} \right) / \left(E_t \left\{ \sum_{j=0}^{\infty} \Pi_{i=0}^j \frac{1 + g_{t+i+1}}{1 + \pi + r_{f,t+i}} \right\} - r_{f,t} \right) \right].$$

Under interesting conditions, such as risk-free rates, equity premiums, and dividend growth rates that conditionally time vary and covary (we consider autoregressive moving average (ARMA) models, autoregressive (AR) models, moving average (MA) models, and correlated errors for dividend growth rates, equity premiums, and interest rates), the individual conditional expectations in equation (8) are analytically intractable.⁷ The difference between the sample mean return and the sample mean risk-free interest rate provides a consistent estimate of π (see Mehra and Prescott (1985)). However, this sample mean difference is imprecisely estimated, even based on more than 100 years of data.

We note that another consistent estimator of π is one that exploits the method of DK (1996). The DK method uses ARMA models for dividend growth rates and interest rates to simulate the conditional expectations

$$E_t \left[\sum_{j=0}^{\infty} \Pi_{i=0}^j \frac{1 + g_{t+i+1}}{1 + \pi + r_{f,t+i}} \right] \quad \text{and} \quad E_{t+1} \left[\sum_{j=0}^{\infty} \Pi_{i=0}^j \frac{1 + g_{t+i+2}}{1 + \pi + r_{f,t+i+1}} \right].$$

⁷Equation (8) is tractable if one assumes joint normality of the logarithms of cash flow growth and discount rates (see Ang and Liu (2001), (2007)). In the general case we consider, the discount rate is a sum of 2 lognormal variables: the interest rate and equity premium. Since such a sum of 2 lognormal variables is not lognormal, our present value discounting problem is not generally nested in the family of models for which analytic solutions are available.

The DK method allows us, for a given *ex ante* equity premium (or a given conditionally varying *ex ante* equity premium process), to simulate the conditional expectations in equation (8) as well as related (unconditional) moments, including the expected dividend yield, return volatility, and *ex post* equity premium. Our estimate of π is produced by finding the value of π that minimizes the distance between the collection of simulated moments (produced by the DK method) and the analogous sample moments (from the U.S. experience over the last half century). The large collection of available moments makes it likely that our analysis can provide a tighter bound on the value of the equity premium than has been achieved previously. The estimation of the conditional expectations in equation (8) relies on the parameters that characterize the dividend growth rate and interest rate models. A joint estimation of these models' parameters and π (i.e., minimizing the distance between simulated and sample moments by varying all of the model's parameters and π at once) would be computationally difficult. We therefore utilize a 2-step procedure in which first, for a given unconditional equity premium, we jointly estimate the parameters that characterize the evolution of dividend growth rates and interest rates. We use these models to simulate data to compare with realized S&P 500 data. Second, we do a grid search over values of the unconditional equity premium to find our SMM estimate of π .

B. The Simulation

Prior to considering the technical details of our simulations, it is useful to consider an overview. Our simulation can be summarized as follows: First we specify assumptions about π , the equity premium anticipated by investors, in particular whether it is constant or conditionally varying and, if it varies over time, how it does so. We also specify the time-series evolution of dividends and interest rates. The choices we make about the processes underlying interest rates, dividend growth rates, and the conditional equity premium define the set of models we consider in the simulation (Models 1–12), and are listed in Table 1. (We explain these choices later.)

The 1st column of Table 1 indicates the numbering that we assign to the models. The 2nd column characterizes the time-series processes used to generate the interest rate, dividend growth rate, and conditional equity premium series. The next 3 columns indicate whether the conditional equity premium process incorporates a downward trend over time (and if so, how much the mean equity premium in 1952 differs from the value in 2004), whether or not there is a structural break (consisting of a 50-bp drop) in the equity premium, and whether or not there is a break in the dividend growth rate process.⁸ The final column in Table 1 indicates which models incorporate uncertainty in generating parameters. Again, we explain all of these features more fully later. In all, we consider 12 representative models, ranging from a simple model with no breaks or trends in the conditional

⁸In each case where we consider model specifications intended to capture real-world features like breaks and trends in rates and premiums, we adopt parameterizations that bias our results to be more conservative (i.e., to produce a wider confidence interval for the equity premium). This allows us to avoid overstating the gains in precision possible with our technique.

TABLE 1
 Characteristics of Simulated Models

In Table 1 we describe the 12 models we consider. The column labeled "Processes for r , g , and π " characterizes the time-series models used to generate the interest rates, dividend growth rates, and equity premium. The column labeled "Downward Trend in Equity Premium Process" identifies whether the conditional equity premium trends downward over the course of the 53-year experiment, and if it does, provides the amount of the downward trend. The next column, "Structural Break in Equity Premium Process" indicates whether the model incorporates a sudden 50-bp drop in the value of the conditional equity premium. The column "Structural Break in Dividend Growth Process" indicates whether the model incorporates a gradual 100-bp increase in the growth rate of the dividend growth rate. The final column indicates that all the models except Models 11 and 12 incorporate sampling variability in generating parameters. Additional model details are as follows. Base Model: Interest rates follow an AR(1), dividend growth rates follow an MA(1), and the equity premium follows an AR(1). Deterministic π Model: Interest rates follow an AR(1), dividend growth rates follow an MA(1), and the equity premium follows a deterministic downward trend with a 50-bp structural break. Best BIC Model:^a Interest rates follow an ARMA(1,1), dividend growth rates follow an MA(1), and the equity premium follows an AR(1). Second-Best BIC Model:^a Interest rates follow an ARMA(1,1), dividend growth rates follow an MA(1), and the equity premium follows an ARMA(1,1).

^aFor Models 7 and 8 we employ the Bayesian information criterion (BIC) to select the order of the ARMA model driving each of the interest rate, equity premium, and dividend growth rate processes. The order of each AR process and each MA process for each series is chosen over a (0, 1, 2) grid. The BIC has been shown by Hannan (1980) to provide consistent estimation of the order of linear ARMA models. We employ the BIC instead of alternative criteria because it delivers relatively parsimonious specifications and because it is widely used in the literature (e.g., Nelson (1991) uses the BIC to select exponential generalized autoregressive heteroskedasticity (EGARCH) models).

Model	Processes for r , g , and π	Downward Trend in Equity Premium Process	Structural Break in Equity Premium Process	Structural Break in Dividend Growth Process	Sampling Variability in Generating Parameters
1	Base model	No	No	No	Yes
2	Base model with π trend	Yes (80 bp)	No	No	Yes
3	Base model with π break	No	Yes (50 bp)	No	Yes
4	Base model with dividend growth trend	No	No	Yes	Yes
5	Base model with π trend and dividend growth trend	Yes (80 bp)	No	Yes	Yes
6	Base model with π break, π trend, and dividend growth trend	Yes (30 bp)	Yes (50 bp)	Yes	Yes
7	Best BIC model ^a with π break, π trend, and dividend growth trend	Yes (30 bp)	Yes (50 bp)	Yes	Yes
8	Second-best BIC model ^a with π break, π trend, and dividend growth trend	Yes (30 bp)	Yes (50 bp)	Yes	Yes
9	Base model with π break and π trend	Yes (30 bp)	Yes (50 bp)	No	Yes
10	Deterministic π model with π break and π trend	Yes (30 bp)	Yes (50 bp)	No	Yes
11	Base model with constant parameters, π break, π trend, and dividend growth trend	Yes (30 bp)	Yes (50 bp)	Yes	No
12	Base model with constant parameters	No	No	No	No

equity premium process (Model 1) to very complex models. After specifying the set of models, we simulate the data and use discounting to estimate stock returns, dividend yields, etc., for each model. Finally, we examine the distributions of the economic variables that emerge from the simulations, and we draw inferences about the equity premium. We now describe each element of our simulation in detail.

1. Processes for the Interest Rate, Dividend Growth Rate, and Equity Premium

The risk-free interest rate and dividend growth rate series we generate are calibrated to the time-series properties of the annually observed 1-year U.S. T-bill

rate and S&P 500 dividend growth rate series over the period from 1952 to 2004.⁹ We considered the ability of various time-series models to eliminate residual autocorrelation and autoregressive conditional heteroskedasticity (ARCH) (evaluated with Lagrange multiplier (LM) tests for residual autocorrelation and for ARCH, both using 5 lags), and we evaluated the log likelihood function and Bayesian information criterion (BIC) across models. Although we will describe the process of model selection 1 variable at a time, our final models were chosen using a full information maximum likelihood (FIML) systems equation estimation and a joint-system BIC optimization.

Economic theory admits a wide range of possible processes for the risk-free interest rate, from constant to autoregressive and highly nonlinear heteroskedastic forms. We find that, in practice, both AR(1) and ARMA(1,1) models of the logarithm of interest rates, based on the model of Hull ((1993), p. 408), perform well in capturing the time-series properties of observed interest rates. We also find that the AR(1) and ARMA(1,1) specifications perform comparably to each other, markedly dominating the performance of other specifications including higher-order models like ARMA(2,2). An attractive feature of modeling the log of interest rates is that doing so restricts nominal interest rates to be positive. Finally, we find standard tests for normality of the error term (and hence conditional lognormality of interest rates) do not reject the null of normality.

Since dividend growth rates have a minimum value of -100% and no theoretical maximum, a natural choice for their distribution is also lognormal. Thus we model the log of 1 plus the dividend growth rate, and we find that both an MA(1) and an AR(1) specification fit the data well, removing evidence of residual autocorrelation and ARCH at 5 lags. These specifications are preferred on the basis of the same criteria used to choose the specification for modeling interest rates. As with the interest rate data, we find standard tests for normality of the error term (and hence conditional lognormality of dividend growth rates) do not reject the null of normality.

Unlike interest rates and dividend growth rates, the ex ante equity premium process is not observable. Nonetheless, the extensively documented predictability of returns (see, for instance Campbell and Shiller (1988) and their references to the literature) strongly suggests that the equity premium varies over time. For our simulations, we model the conditional equity premium process based on observable factors, exploiting Merton's (1980) conditional capital asset pricing model (CAPM). The fitted equity premium value from Merton's conditional CAPM is modeled as a time-series process, covarying with interest rates and dividend growth rates. This system of variables (i.e., dividend growth rates, interest rates,

⁹Our findings with respect to the value and precision of the equity premium estimate are not sensitive to the use of quarterly versus annual data, though differences do arise in the time-series models for dividend growth rates and interest rates. With quarterly data, dividend growth rates follow an ARMA(1,1) process, and quarterly dividends show more persistence than annual dividends. Similarly, quarterly interest rates require an ARMA(1,1) model to remove evidence of residual autocorrelation, whereas a lower-order model suffices with annual data. The explanatory power of these models is higher with quarterly data, and the magnitude of coefficients capturing the autocorrelation, such as the AR(1) term, is larger with quarterly data.

and equity premiums) is then simulated, and these simulated values are used to calculate prices and returns, as described later.

Merton's (1980) conditional CAPM is expressed in terms of returns in excess of the risk-free rate, or in other words, the period-by-period equity premium. For the market return, Merton derives the expression

$$(9) \quad E_t [\tilde{R}_{m,t+1}] = \lambda \text{var}_t(\tilde{R}_{m,t+1}),$$

where $\tilde{R}_{m,t+1}$ is the excess return on the market portfolio and $\text{var}_t(\tilde{R}_{m,t+1})$ is the time-varying conditional variance of the market excess return. Merton argues that λ in equation (9) is the weighted sum of the reciprocal of each investor's coefficient of relative risk aversion, with the weight being related to the distribution of wealth among individuals. Equation (9) defines a conditional equity premium but has the equity premium varying only as a function of the conditional variance. Following Bekaert and Harvey (1995), it is possible to allow λ to vary over time by making it a parametric function of conditioning variables (indicated below as Z_t). The functional form Bekaert and Harvey employ (in equation (12) of their paper) is exponential, restricting the price of risk to be positive: $\lambda_t = \exp(\delta'Z_t)$.

Shiller (1984), Rozeff (1984), Campbell and Shiller (1988), Hodrick (1992), and Bekaert and Harvey (1995) all document the usefulness of dividend yields to predict returns, so we use lagged dividend yields as our conditioning variable. We make use of a simple ARCH specification to model $\text{var}_t(\tilde{R}_{m,t+1})$. We estimate the following model (calculating \tilde{R} by making use of annual S&P 500 returns and 1-year U.S. Treasury rates from 1952 to 2004):

$$(10) \quad \tilde{R}_{m,t+1} = \lambda_t \text{var}_t(\tilde{R}_{m,t+1}) + e_{m,t+1},$$

$$(11) \quad \text{var}_t(\tilde{R}_{m,t+1}) = \omega + \alpha e_{m,t}^2,$$

$$(12) \quad \lambda_t = \exp\left(\delta_0 + \delta_1 \frac{D_t}{P_t}\right).$$

The values of estimated parameters are (with standard errors in parentheses) $\delta_0 = -3.93$ (0.646), $\delta_1 = 0.277$ (0.118), $\omega = 0.0194$ (0.0033), and $\alpha = 0.542$ (0.346). The R^2 of this model is 2.8%.

For our simulations, we model the time-series process of the conditional equity premium (denoted π_t) by using the excess return as a proxy for the equity premium: $\hat{\pi}_{t+1} = \hat{\lambda}_t \hat{\text{var}}_t(\tilde{R}_{m,t+1})$, where $\hat{\lambda}_t = \exp(-3.93 + 0.277D_t/P_t)$, $\hat{\text{var}}_t(\tilde{R}_{m,t+1}) = 0.0194 + 0.542\hat{e}_{m,t}^2$, and $\hat{e}_{m,t} = \tilde{R}_{m,t} - \hat{\pi}_t$. The conditional equity premium we estimate here, $\hat{\pi}_t$, follows a strong AR(1) time-series process, similar to that of the risk-free interest rate,¹⁰ so that when the equity premium is perturbed it reverts to its mean slowly. The persistence permits slightly more volatile returns in our simulations than would otherwise be the case. The autocorrelation is driven in part by the persistence of D/P and in part by the persistence of the conditional variance. We find that standard tests for normality of the error term

¹⁰The mean of the estimated equity premium from this model is 5.8% and its standard deviation is 2.2%. An AR(1) model of the natural logarithm of the equity premium, $\hat{\pi}_t$, has a coefficient of 0.79 on the lagged equity premium, with a standard error of 0.050 and an R^2 of 0.83.

(and hence conditional lognormality of the equity premium) show some evidence of nonnormality when estimated as a single equation, but less or no evidence of nonnormality if estimated in a system of equations that includes the interest rate and dividend growth rate equations.

We generate the conditional equity premiums, interest rate, and dividend growth rate series as autocorrelated series with jointly normal error terms, calibrated to the degree of autocorrelation observed in the U.S. data. The processes we simulate also mimic the covariance structure between the residuals from the time-series models of equity premiums, interest rates, and dividend growth rates as estimated using U.S. data. We adjust the mean and standard deviation of these lognormal processes to generate the desired level and variability for each when they are transformed back into levels. The coefficients and error covariance structure are estimated with FIML. Very similar results are obtained using iterative GMM and Newey and West (1987) heteroskedasticity and autocorrelation consistent covariance estimation.

To give a sense for what our estimated models for interest rates, dividend growth rates, and the equity premium look like, we present in Table 2 the estimated parameters of Model 1, which incorporates an AR(1) model for interest rates (r), an MA(1) model for dividend growth rates (g), and an AR(1) model for the equity premium (π).

TABLE 2
Estimated Parameters of Model 1

In Table 2 we present parameter estimates for Model 1, with standard errors of the estimates in parentheses. The covariance of $\epsilon_{r,t}$ and $\epsilon_{g,t}$ equals 0.00240, the covariance of $\epsilon_{r,t}$ and $\epsilon_{\pi,t}$ equals -0.0117, and the covariance of $\epsilon_{g,t}$ and $\epsilon_{\pi,t}$ equals 0.0018. The variance of $\epsilon_{r,t}$ equals 0.0890, the variance of $\epsilon_{g,t}$ equals 0.000986, and the variance of $\epsilon_{\pi,t}$ equals 0.0648. The adjusted R^2 for the interest rate equation is 72.9%, the adjusted R^2 for the dividend growth rate equation is 30.0%, and the adjusted R^2 for the equity premium equation is 79.5%.

Dependent Variable		Intercept (Standard Error)	Explanatory Variable (Standard Error)	Residual
$\log(r_{t,t})$	=	-0.214 (0.262)	+ 0.929 $\log(r_{t,t-1})$ (0.086)	+ $\epsilon_{r,t}$
$\log(1 + g_t)$	=	0.0516 (0.0063)	+ 0.454 $\epsilon_{g,t-1}$ (0.084)	+ $\epsilon_{g,t}$
$\log(\hat{\pi}_t)$	=	-0.562 (0.230)	+ 0.851 $\log(\hat{\pi}_{t-1})$ (0.070)	+ $\epsilon_{\pi,t}$

2. Allowing a Downward Trend in the Equity Premium Process

Pástor and Stambaugh (2001), among others, provide evidence that the equity premium has been trending downward over the sample period we study, finding a modest downward trend of roughly 0.8% in total since the early 1950s, with much of the difference coming from a steep decline in the 1990s. Their study of the equity premium has the premium fluctuating between about 4% and 6% since 1834. Given this evidence and the fact that we calibrate to data starting in the 1950s, some of the specifications we investigate later permit a 0.8% trend in the equity premium. When modeling a trend with a break, we limit ourselves to a 0.3% trend with an additional 0.5% break, as discussed later. This is accomplished in conjunction with setting the equity premium to follow an AR(1) process.

3. Allowing a Structural Break in the Equity Premium Process

Pástor and Stambaugh (2001) estimate the probability of a structural break in the equity premium over the last 2 centuries. They find fairly strong support for there having been a structural break over the 1990s that led to a 0.5% drop in the equity premium. An aggressive interpretation of their results would have the majority of the drop in the equity premium over the 1990s occurring at once. We decide to adopt a one-time-drop specification because doing so makes our results more conservative (i.e., produces a wider confidence interval for the unconditional equity premium). Spreading the drop in the premium across several years serves only to narrow the range of the equity premium consistent with the U.S. returns data over the last 50 years, which would only bolster our claims to provide a much tighter confidence interval about the estimate of the equity premium. In some of the specifications we investigate later we therefore incorporate an abrupt 0.5% drop in the equity premium. (We time this drop to coincide with 1990, 39 years into our simulation period.) This feature of the equity premium process can be incorporated in the models we use with or without incorporating the other features discussed previously; as discussed later, we investigate combinations and permutations of these various features.

4. Allowing for Sampling Variability in Generating Parameters

In order to produce our simulations we must first estimate parameter values for the time-series models of dividend growth rates, interest rates, and conditional equity premiums. These estimates themselves incorporate sampling variability. Fortunately, estimates of the sampling variability are available to us through the covariance matrix of our parameters, so we can incorporate uncertainty about the true values of these parameters into our simulations. We estimate our system of equations (the dividend growth rate, interest rate, and conditional equity premium) jointly with FIML, and generate for *each* simulation an independent set of parameters drawn randomly from the joint limiting normal distribution of these parameter estimates (including the variance and covariance of the equation residuals) subject to some technical considerations¹¹ and data consistency checks.¹² This process accounts for possible variability in the true state of the world that generates dividends, interest rates, and equity premiums.

To illustrate, consider Model 1 (recall that Table 2 contains the parameter estimates associated with Model 1):

$$\begin{aligned} \log(r_{f,t}) &= \alpha_r + \rho_r \log(r_{f,t-1}) + \epsilon_{r,t}, \\ \log(1 + g_t) &= \alpha_g + \theta_g \epsilon_{g,t-1} + \epsilon_{g,t}, \quad \text{and} \\ \log(\hat{\pi}_t) &= \alpha_\pi + \rho_\pi \log(\hat{\pi}_{t-1}) + \epsilon_{\pi,t}. \end{aligned}$$

¹¹The time-series models must exhibit stationarity, the growth rate of dividends must be strictly less than the discount rate, and the residual variances must be greater than 0.

¹²The parameters must generate mean interest rates, dividend growth rates, and ex post equity premiums that lie within 3 standard deviations of the U.S. data sample mean. Also, the limiting price-dividend ratio must be within 50 standard deviations of the mean U.S. price-dividend ratio. This last consistency check rules out some extreme simulations generated when the random draw of parameters leads to near unit root behavior. The vast majority of simulations do not exhibit price-dividend ratios that are more than a few standard deviations from the mean of the U.S. data.

We provide the estimated covariance matrix of the *parameter estimates* for this model in Table 3. The estimated covariance matrix of the *equation residual variances* is shown in Table 4. (The variances of the residuals from Model 1 are reported in the notes to Table 2.)

TABLE 3
Estimated Covariance Matrix for Model 1 Parameters

In Table 3 we present the variance-covariance matrix for the Model 1 parameter estimates. Model 1 is generated using these processes for the risk-free rate, the dividend growth rate, and the conditional equity premium:

$$\begin{aligned} \log(r_{f,t}) &= \alpha_r + \rho_r \log(r_{f,t-1}) + \epsilon_{r,t}, \\ \log(1 + g_t) &= \alpha_g + \theta_g \epsilon_{g,t-1} + \epsilon_{g,t}, \text{ and} \\ \log(\hat{\pi}_t) &= \alpha_\pi + \rho_\pi \log(\hat{\pi}_{t-1}) + \epsilon_{\pi,t}. \end{aligned}$$

The top-left element in the table, equal to 0.068705, is the variance of the parameter estimate of α_r . The entry below the top-left element, equal to 0.022307, is the covariance between the estimate of α_r and ρ_r , and so on.

Parameter	α_r	ρ_r	α_g	θ_g	α_π	ρ_π
α_r	0.068705	0.022307	-0.000051933	0.000226443	-0.012165	-0.003511
ρ_r	0.022307	0.007436	-0.000040346	0.000114831	-0.004730	-0.001401
α_g	-0.000052	-0.000040	0.000039674	0.000025651	0.000153	0.000031
θ_g	0.000226	0.000115	0.000025651	0.007086714	0.001699	0.000454
α_π	-0.012165	-0.004730	0.000153376	0.001699151	0.052664	0.015791
ρ_π	-0.003511	-0.001401	0.000031495	0.000453874	0.015791	0.004844

TABLE 4
Estimated Covariance Matrix of Model 1 Residual Variances

In Table 4 we present the estimated covariance matrix for the Model 1 residual variances. Model 1 is generated using these processes for the risk-free rate, the dividend growth rate, and the conditional equity premium:

$$\begin{aligned} \log(r_{f,t}) &= \alpha_r + \rho_r \log(r_{f,t-1}) + \epsilon_{r,t}, \\ \log(1 + g_t) &= \alpha_g + \theta_g \epsilon_{g,t-1} + \epsilon_{g,t}, \text{ and} \\ \log(\hat{\pi}_t) &= \alpha_\pi + \rho_\pi \log(\hat{\pi}_{t-1}) + \epsilon_{\pi,t}. \end{aligned}$$

The top-left element in the table, equal to 0.0000944, is the variance of ϵ_r^2 . The entry below the top-left element, equal to 1.9729×10^{-6} , is the covariance between the estimate of ϵ_r^2 and the product of ϵ_r and ϵ_g , and so on.

Parameter	ϵ_r^2	$\epsilon_r \epsilon_g$	$\epsilon_r \epsilon_\pi$	ϵ_g^2	$\epsilon_g \epsilon_\pi$	ϵ_π^2
ϵ_r^2	0.0000944	1.9729×10^{-6}	-8.351×10^{-7}	-1.902×10^{-7}	-1.564×10^{-6}	-1.69×10^6
$\epsilon_r \epsilon_g$	1.9729×10^{-6}	8.5163×10^{-7}	1.0437×10^{-6}	4.3066×10^{-8}	-1.602×10^{-7}	9.1448×10^{-7}
$\epsilon_r \epsilon_\pi$	-8.351×10^{-7}	1.0437×10^{-6}	0.0000797	1.8827×10^{-7}	5.001×10^{-6}	-0.000044
ϵ_g^2	-1.902×10^{-7}	4.3066×10^{-8}	1.8827×10^{-7}	4.8337×10^{-8}	9.6885×10^{-8}	1.3458×10^{-6}
$\epsilon_g \epsilon_\pi$	-1.564×10^{-6}	-1.602×10^{-7}	5.001×10^{-6}	9.6885×10^{-8}	3.5567×10^{-6}	0.0000203
ϵ_π^2	-1.69×10^{-6}	9.1448×10^{-7}	-0.000044	1.3458×10^{-6}	0.0000203	0.0005009

Exploiting block diagonality of the parameters of the mean and variance, and asymptotic normality of all the estimated parameters, we generate 2 sets of normally distributed random variables. Each set is independent of the other, the 1st set of 6 having the covariance matrix from Table 3 with means equal to the parameter estimates listed in Table 2, and the 2nd set of 6 having the covariance matrix from Table 4, with means equal to the equation residual covariances listed in the notes to Table 2. This set of 12 random variables is then used to simulate interest rates, dividend growth rates, and equity premiums, subject to the consistency checks footnoted earlier.

5. Allowing for Disappearing Dividends

Since the late 1970s, U.S. firms have been increasingly delivering cash flows to investors via share repurchases and other means that are not officially recorded as “corporate dividends.” Fama and French (2001) document this widespread decline of regular dividend payments starting in 1978, consistent with evidence provided by Bagwell and Shoven (1989) and others. However, Fama and French also find evidence that the disappearance of dividends is due in part to an increase in the inflow of new listing to U.S. stock exchanges, representing mostly young companies that would not be expected to pay dividends. Further, they find that there is a much smaller decline in the propensity to pay dividends among the type of firms that we calibrate to, those in the S&P 500 index.

Consistent with Fama and French (2001), we find no evidence of a break in our data on dividend growth rates. Although dividend *yields* on the S&P 500 index have dropped dramatically over time, dividend growth rates have not. The decline in yields has been a function of prices rising faster than dividends since 1978, not a slowdown in dividend growth. From 1952 through 1978, the year Fama and French document as the year of the structural break in dividend payments, dividend growth rates among the S&P 500 firms averaged 4.9% with an annual standard deviation of 3.9%, and from 1979 to 2000 the dividend growth rates averaged 5.5% with an annual standard deviation of 3.8%, virtually indistinguishable from the pre-1979 period. Time-series properties pre- and post-1978 are also very similar across these 2 periods. Consistent with this stability of dividend growth pre- and post-1978 and Bagwell and Shoven’s (1989) documentation of increased share repurchases in the 1980s, *earnings growth rates* of firms in the S&P 500 index have accelerated since the 1952–1978 period, from 6.8% pre-1979 to 7.8% post-1978. Similar to the dividend growth rate data, the time-series properties of the earnings growth rate data did not change.

In order to determine the sensitivity of our experiments to mismeasurement of cash flows to investors, we consider a dividend growth rate process with a structural break 27 years into the time series to correspond to a possible break in our dividend data for the S&P 500 data after 1978. We calibrate to the S&P 500 earnings data mean growth rate increase over 1979–2000, an upward shift of 100 bp, to proxy for the increase in total cash flows to investors. That is, we increase the growth rate of dividends by 5 bp a year for 20 years, starting in year 27 of the simulation (corresponding to 1978 for the S&P 500 data), to increase the mean growth rate of our dividend growth series 100 bp, mimicking the proportional increase in earnings growth rates observed in the data.

6. Numerical Simulation Details

In Section III we report results from our models. Before doing so, however, it is useful to outline the steps used in our simulations. We therefore now illustrate for the n th simulated economy how prices (P_t^n), returns (R_t^n), ex post equity premiums ($\hat{\pi}^n$), etc. (where $n = 1, \dots, N$ and $t = 1, \dots, T$) are produced given assumed processes for dividends, dividend growth rates, risk-free interest rates, and the

(ex ante) equity premium of the n th economy.¹³ For the purpose of this illustration, we assume fixed parameters (i.e., no parameter uncertainty), a constant unconditional equity premium, and an AR(1) model for interest rates. Further, to illustrate the procedure required for a moving average error model, we assume an MA(1) process for dividend growth rates. Relaxing these assumptions (incorporating parameter uncertainty, ARMA(1,1) processes for interest rates and dividend growth rates, and a conditionally varying equity premium) complicates the procedure outlined later only slightly. We simulate the economies out for 50 periods, then at period 51 we start our calculation of market prices, returns, etc. (to avoid contaminating the simulations with the initial conditions). For simplicity, we omit this detail from the description later.


Our fundamental pricing relationship, equation (5), rewritten to correspond to the n th economy, is

$$(13) \quad P_t^n = D_t^n E_t \left\{ \sum_{i=0}^{\infty} \left(\Pi_{k=0}^i \left[\frac{1 + g_{t+k+1}^n}{1 + r_{t+k}^n} \right] \right) \right\}.$$

Note that $r_{t+k}^n = r_{f,t+k}^n + \pi$ where π is the ex ante equity premium. We produce prices (P_t^n), returns (R_t^n), and ex post equity premiums ($\hat{\pi}^n$), utilizing only dividend growth rates (g_{t+k+1}^n) and discount rates (r_{t+k}^n). Figure 1 summarizes the 3 steps of the Monte Carlo simulation.

FIGURE 1
Diagram of a Simple Market Price Calculation
for the t th Observation of the n th Economy

Figure 1 depicts the evolution of interest rates and dividend growth rates upon which the simulations are based.



Step 1. In forming P_t^n , the most recent fundamental information available to an investor would be g_t^n , D_t^n , and r_t^n . Thus g_t^n , D_t^n , and r_t^n must be generated directly in our simulations, whereas P_t^n is formed by numerical integration based on the evolution of g , D , and r . The objective of Steps 1a–1c outlined later is to produce dividend growth and interest rates that replicate real-world dividend growth and interest rate data. That is, the simulated dividend growth and interest rates must have the same mean, variance, covariance, and autocorrelation structure as observed S&P 500 dividend growth rates and U.S. interest rates. In terms of Figure 1, Step 1 forms g_t^n , D_t^n , and r_t^n only.

¹³We set the number of economies, N , at 2,000. This is a sufficiently large number of replications to produce results with very small simulation error.

Step 1a. Set the initial dividend, D_0^n , equal to 1, and the lagged innovation of the logarithm of the dividend growth rates, $\epsilon_{g,0}^n$, to 0. To match the real-world interest rate data, set $\log(r_{f,0}^n) = -2.90$ (the mean value of log interest rates). Generate 2 independent standard normal random numbers, η_1^n and ν_1^n (note that the subscript on these random numbers indicates time, $t=1$), and form 2 correlated random variables, $\epsilon_{r,1}^n = 0.319(0.25\eta_1^n + (1 - 0.25^2)^{0.5}\nu_1^n)$ and $\epsilon_{g,1}^n = 0.0311\eta_1^n$. These are the simulated innovations to the interest rate and dividend growth rate processes, formed to have standard deviations of 0.319 and 0.0311, respectively, to match the data and to be correlated with correlation coefficient 0.25 as we find in the S&P 500 return and T-bill rate data. Next, form $\log(1 + g_1^n) = 0.049 + 0.64\epsilon_{g,0}^n + \epsilon_{g,1}^n$ and $\log(r_{f,1}^n) = -0.35 + 0.88 \log(r_{f,0}^n) + \epsilon_{r,1}^n$ to match the parameters estimated on the S&P 500 index data 1952–2004 of these models.¹⁴ Also form $D_1^n = D_0^n(1 + g_1^n)$.

Step 1b. Produce 2 correlated normal random variables, $\epsilon_{r,2}^n$ and $\epsilon_{g,2}^n$, as in Step 1a, and conditioning on $\epsilon_{g,1}^n$ and $\log(r_{f,1}^n)$ from Step 1a produce $\log(1 + g_2^n) = 0.049 + 0.64\epsilon_{g,1}^n + \epsilon_{g,2}^n$, $\log(r_{f,2}^n) = -0.35 + 0.88 \log(r_{f,1}^n) + \epsilon_{r,2}^n$, and $D_2^n = D_1^n(1 + g_2^n)$.

Step 1c. Repeat Step 1b to form $\log(1 + g_t^n)$, $\log(r_{f,t}^n)$, and D_t^n for $t = 3, 4, 5, \dots, T$ and for each economy $n=1, 2, 3, \dots, N$. Then calculate the dividend growth rate g_t^n and the discount rate r_t^n .

Step 2. For each time period $t=1, 2, 3, \dots, T$ and economy $n=1, 2, 3, \dots, N$, we calculate prices, P_t^n . In order to do this we must solve for the expectation of the infinite sum of discounted future dividends conditional on time t information for economy n . That is, we must produce a set of possible paths of dividends and interest rates that might be observed in periods $t+1, t+2, \dots$, given what is known at period t , and use these to solve the expectation of equation (13). We use the superscript j to index the possible paths of future economies that could evolve from the current state of the economy. In Step 2d, we describe how we are able to solve for the expectation of an infinite sum using a *finite* stream of future dividends.

Step 2a. Set $\epsilon_{g,t}^{j,n} = \epsilon_{g,t}^n$ and $\log(r_{f,t}^{j,n}) = \log(r_{f,t}^n)$ for $j = 1, 2, 3, \dots, J$.¹⁵ Generate 2 independent standard normal random numbers, $\eta_{t+1}^{j,n}$ and $\nu_{t+1}^{j,n}$, and form 2 correlated random variables $\epsilon_{r,t+1}^{j,n} = 0.319(0.25\eta_{t+1}^{j,n} + (1 - 0.25^2)^{0.5}\nu_{t+1}^{j,n})$ and

¹⁴Note that by construction these parameters do not match those reported in Table 2, as this system does not incorporate a conditionally varying equity premium.

¹⁵We choose J to lie between 1,000 and 100,000, to ensure the Monte Carlo simulation error in calculating prices and returns is controlled to be less than 0.20%. For the typical case the simulation error is far less than 0.20%. To determine the simulation error, we conducted a simulation of the simulations. Unlike some Monte Carlo experiments (such as those estimating the size of a test statistic under the null), the standard error of the simulation error for most of our estimates (returns, prices, etc.) are themselves analytically intractable and must be simulated. In order to estimate the standard error of the simulation error in estimating market prices, we estimated a single market price 2,000 times, each time independent of the other, and from this set of prices compute the mean and variance of the price estimate. With the number of possible paths, J , equal to no less than 1,000, we find that the standard deviation of the simulation error is less than 0.20% of the price, which is sufficiently small as not to be a source of concern for our study. The number of simulations has to be substantially greater than 1,000 for some cases depending on the model specification.

$\epsilon_{g,t+1}^{j,n} = 0.0311\eta_{t+1}^{j,n}$ for $j = 1, 2, 3, \dots, J$.¹⁶ These are the simulated innovations to the interest rate and dividend growth rate processes, respectively. Form $\log(1 + g_{t+1}^{j,n}) = 0.049 + 0.64\epsilon_{g,t}^{j,n} + \epsilon_{g,t+1}^{j,n}$ and $\log(r_{f,t+1}^{j,n}) = -0.35 + 0.88 \log(r_{f,t}^{j,n}) + \epsilon_{r,t+1}^{j,n}$.

Step 2b. Produce 2 correlated normal random variables $\epsilon_{r,t+2}^{j,n}$ and $\epsilon_{g,t+2}^{j,n}$ as in Step 2a, and conditioning on $\epsilon_{g,t+1}^{j,n}$ and $\log(r_{f,t+1}^{j,n})$ from Step 2a produce $\log(1 + g_{t+2}^{j,n}) = 0.049 + 0.64\epsilon_{g,t+1}^{j,n} + \epsilon_{g,t+2}^{j,n}$ and $\log(r_{f,t+2}^{j,n}) = -0.35 + 0.88 \log(r_{f,t+1}^{j,n}) + \epsilon_{r,t+2}^{j,n}$ for $j = 1, 2, 3, \dots, J$.

Step 2c. Repeat Step 2b to form $\log(1 + g_{t+i}^{j,n})$ and $\log(r_{f,t+i}^{j,n})$ for $i = 3, 4, \dots, I$, $j = 1, 2, 3, \dots, J$, and economies $n = 1, 2, 3, \dots, N$.

Step 2d. The discounted present value of each of the individual J streams of dividends is now taken in accordance with equation (13), with the j th present value price noted as $P_t^{j,n}$. Finally, the price for the n th economy in period t is formed: $P_t^n = \frac{1}{J} \sum_{j=1}^J P_t^{j,n}$. In considering these prices, note that according to equation (13) the value of I in Step 2c above should run out to infinity. Since the ratio of gross dividend growth rates to gross discount rates is less than unity in steady state, the individual product elements in the infinite sum in equation (13) eventually converge to 0 as I increases. We find $I = 1,000$ is large enough in our simulations so that the truncation does not materially affect our results.

Step 3. After performing Steps 1a–1c and 2a–2d for $t = 1, \dots, T$, rolling out N independent economies for T periods, we construct the market returns for each economy, $R_{t+1}^n = (P_{t+1}^n + D_{t+1}^n - P_t^n) / P_t^n$, and the ex post equity premium that agents in the n th economy would observe, $\hat{\pi}^n$, estimated from equation (1) as the mean difference in market returns and the risk-free rate. The length of the time series T is chosen to be 53 to imitate the 53 years of annual data we have available for the S&P 500 from 1952 to 2004.

III. Univariate Conditional Distributions for Model 1

All of the results in this section of the paper are based on Model 1, as defined in Table 1. Model 1 incorporates interest rates that follow an AR(1) process and dividend growth rates that follow an MA(1) process. The conditional equity premium in Model 1 follows an AR(1) process (that emerges from the fitted equity premium produced by Merton's (1980) conditional CAPM, as detailed in Section II.B.1), with no trends or breaks in either the equity premium process or dividend growth rate process. We start with this "plain vanilla" model because it provides a good illustration of how well dividend-discounting models that incorporate autocorrelated dividend growth and discount rate processes can produce prices and returns that fit the experience of the last half century in the U.S. This model also provides a good starting point to contrast with models employing breaks and trends in equity premium and dividend growth processes. We consider

¹⁶For our random number generation we make use of a variance reduction technique, stratified sampling. This technique has us drawing pseudo-random numbers ensuring that $q\%$ of these draws come from the q th percentile, so that our sampling does not weight any grouping of random draws too heavily.

more complex and arguably more realistic models incorporating trends and breaks later in the paper.

We first consider what we can learn by looking at the univariate statistics that emerge from our simulations. We can use the univariate distributions to place loose bounds on plausible values of the equity premium. While the analysis in this section based on univariate empirical distributions is somewhat casual, in Section IV we conduct formal analysis based on χ^2 statistics and the joint distributions of the data, yielding very tight bounds on plausible values of the equity premium and identifying plausible models of the equity premium process.

We can compare any financial statistic from the last half century to our simulated economies provided the statistic is based on returns, dividends, or prices, as these are data that the simulation produces. We could also consider moments based on interest rates or dividend growth rates, but since we calibrate our models to interest rates and dividend growth rates, all our simulations should (and do) fit these moments well by construction. We choose moments based on 2 considerations. First, the moments should be familiar and the significance of the moments to economic theory should be obvious. Second, the moments should be precisely estimated; if the moments are too “noisy,” they will not help us narrow the range of equity premiums (e.g., return skew and kurtosis are very imprecisely estimated with even 50 years of data, so that these moments are largely uninformative). The moments must also be well defined (e.g., moments must be finite). The expected value of the price of equity is undefined, but we can use prices in concert with a cointegrated variable like lagged price (to form returns) or dividends (to form dividend yields).


We summarize the simulation results with plots of probability distributions of the simulated data for various financial statistics. Graphs A–D of Figure 2 contain probability distribution functions (PDFs) corresponding to the mean ex post equity premium, the mean dividend yield, the Sharpe ratio, and return volatility respectively, based on assumed unconditional equity premiums of 2.75%, 3.75%, 5%, and 8%. The dotted lines depicting the PDFs in Figure 2 are thinnest for the 2.75% case and become progressively thicker for the 3.75%, 5%, and 8% cases. The realized U.S. data are denoted in each panel with a solid vertical line. Comparison of the simulated distribution with realized values in these plots permits a very quick, if casual, first assessment of how well the realized U.S. data agree with the simulated data.

The realized U.S. mean excess return, displayed as the vertical line in Graph A of Figure 2, is furthest in the right tail of the distribution corresponding to a 2.75% unconditional equity premium, and furthest in the left tail for the equity premium of 8%. The wide range of the distribution of the mean excess return for each assumed value of the unconditional equity premium is consistent with the experience of the last half century in the U.S., in which the mean excess return has a 95% confidence interval spanning plus or minus roughly 4% or 5%. The realized dividend yield of 3.4%, displayed in Graph B, is unusually low for the 5% and 8% equity premium cases, but it is near the center of the distribution for the equity premium values of 2.75% and 3.75%. In Graph C, only the Sharpe ratios generated with an equity premium of 8% appear inconsistent with the U.S. experience of the last half century. The return volatility, displayed

FIGURE 2

Probability Distribution Functions of Simulated Ex Post Equity Premiums, Dividend Yields, Sharpe Ratios, and Return Standard Deviations

Figure 2 contains probability distribution functions (PDFs) for various financial statistics generated in 2,000 simulated economies based on Model 1 from Table 1. Each graph contains a PDF for each of 4 different assumed values of the unconditional equity premium: 2.75%, 3.75%, 5%, and 8%. Graph A shows the distribution of the ex post equity premium (mean return minus mean interest rate), Graph B shows the mean dividend yield distribution (dividend divided by price), Graph C shows the Sharpe ratio distribution (excess return divided by the standard deviation of the excess return), and Graph D shows the distribution of the standard deviation of excess returns. In each graph, a vertical line indicates the U.S. data realized over 1952–2004, the value of the estimated ex post equity premium, mean dividend yield, mean Sharpe ratio, and excess return standard deviation, respectively. The simulated statistics are estimated on 53 years of generated data for each economy, mimicking the data period we used to estimate the realized U.S. results.



in Graph D, clearly indicates that the experience of the U.S. over the last half century is somewhat unusual for all values of the equity premiums considered, though least unusual for the lowest equity premium.

Figure 2 has 2 central implications of interest to us. First, the financial variable statistics produced in our simulations are broadly consistent with what has been observed in the U.S. economy over the past 5 decades. Most simulated statistics, including those based on prices and returns, match the magnitudes of financial quantities from the realized U.S. data, even though we do not calibrate our models to prices or returns. Second, the results suggest that the 2.75%–8% interval we present here likely contains the equity premium consistent with the U.S. economy.

Univariate results for Models 2–10 are qualitatively very similar to those presented for Model 1. Univariate results for Models 11 and 12, in contrast, are grossly rejected by the experience of the U.S. economy. Detailed univariate results for Models 2–12 are omitted for the sake of brevity, but the poor performance of Models 11 and 12 will be evident in multivariate results reported later.

IV. Model Extensions, Multivariate Analysis, and Tests

To narrow further the range of plausible equity premium values, we exploit the full power of our simulation procedure by considering the *joint distributions* of statistics that arise in our simulations and comparing them to empirical moments of the data. The 3 moments of the data that we focus on (the mean ex post equity premium, excess return volatility, and mean dividend yield) have the advantage of being the most precisely estimated and hence the most informative for the value of the equity premium.

Our purpose in considering joint distributions is twofold. First, multivariate tests are used to form a tight confidence bound on the equity premium. These tests strongly reject our models if the unconditional equity premium is outside of a narrow range around 3.5%. This range is not sensitive to even fairly substantial changes in the model specification, which suggests that the 3.5% finding is robust. Second, the analysis exploiting the joint distribution leads to rejection of model specifications that fail to incorporate certain features, such as trends and breaks in the equity premium. Interestingly, even when a model specification is rejected, we find the most plausible equity premium still lies in the same range as the rest of our models, very near 3.5%.

Up to this point we have considered detailed results for Model 1 exclusively, but Model 1 does not incorporate a gradual downward trend in the equity premium, a structural break in the equity premium process, or an increase in the growth rate of cash flows. Thus we now consider models that incorporate 1, 2, or all 3 of these features, as well as different time-series models for interest rates and equity premiums. We also consider stripped-down models to assess the marginal contribution of model features such as parameter uncertainty and the specification of the time-series process used to model dividend growth rates and interest rates.

In Figures 3–5 (to be fully discussed later), we present χ^2 test statistics for the null hypothesis, that the U.S. experience during 1952–2004 could have been a random draw from the simulated distribution of the mean ex post equity

premium, the excess return volatility, and the mean dividend yield.¹⁷ A significant test statistic, in this context, suggests that the combination of financial statistics observed for the U.S. economy is significantly unusual compared to the collection of simulated data, leading us to reject the null hypothesis that the given model and assumed equity premium value could have generated the U.S. data of the last half century. It is possible to reject every equity premium value if we use models of the equity premium that are misspecified (the rejection of the null hypothesis can be interpreted as a rejection of the model). It is also possible that a very wide range of equity premium values are not rejected for a collection of models, thwarting our efforts to provide a precise estimate of the equity premium or a small range of allowable equity premium models.

Our results reveal that models that ignore breaks and trends in the conditional equity premium process are rejected for virtually every value of the unconditional equity premium we consider. For only a group of sophisticated models that incorporate trends and breaks in the equity premium do we fail to reject a narrow range of unconditional equity premiums, roughly between 3% and 4%. We tend to reject models if the impact on cash flows to shareholders from share repurchases are ignored. We begin with some simple models, then consider models that are arguably more realistic, as they incorporate equity premium and cash flow trends and breaks, and we finish by considering a host of related issues, including the impact of parameter estimation error and, separately, *investor* uncertainty about the fundamental value of equities.

A. Simple (One-at-a-Time) Model Extensions

We now consider extensions to Model 1, each extension adding a single feature to the base model. Recall that the features of each model are summarized in Table 1. For Model 2, an 80-bp downward trend is incorporated in the equity premium process. For Model 3, a 50-bp drop in year 39 of the simulation (corresponding to 1990 for the S&P 500 data) is incorporated in the equity premium process. For Model 4, the dividend growth rate process is shifted gradually upward a total of 100 bp, starting in year 27 of the simulation (corresponding to 1978 for the S&P 500 data) and continuing for 20 years at a rate of 5 bp per year. These models help us evaluate if one or another feature documented in the literature can markedly improve model performance over the simple base model.

Graph A of Figure 3 displays plots of the value of joint χ^2 tests on 3 moments of the data, the mean ex post equity premium, the excess return volatility, and the

¹⁷The χ^2 tests are based on joint normality of sample estimates of moments of the simulated data, which follow an asymptotic normal distribution based on a law of large numbers (see White (1984) for details). In the case of excess return volatility, we consider the *cube root* of the return variance, which is approximately normally distributed (see p. 399 of Kendall and Stuart (1977) for further details). We also estimate the probability of rejection using bootstrapped *p*-values, to guard against deviations from normality. These bootstrapped values are qualitatively identical to the asymptotic distribution *p*-values. Finally, when performing tests that include the dividend yield moment, if the simulation includes a break in dividends corresponding to an increase in cash payouts starting in 1978 in the U.S. data (again, see Fama and French (2001)), we also adjust the U.S. data to reflect the increase in mean payout levels. This makes for a small difference in the mean U.S. payout ratio and no qualitative change to our results if ignored.

mean dividend yield, for Models 1–4, and shows how the test statistic varies as the unconditional equity premium varies from 2.25% to 8% in increments as small as $\frac{1}{8}$ of a percent toward the lower end of that range. Graphs B–D of Figure 3 display the univariate student t -test statistics for each of these 3 moments of the data, again showing how the test statistic varies with the assumed value of the equity premium. The equity premium values indicated on the horizontal axis represent the *ending* values of the equity premium in each set of simulations. For models that incorporate a downward trend or a structural break in the equity premium process, the ending value of the equity premium differs from the starting value. So, for instance, Model 2 has a starting equity premium that is 80 bp higher than that displayed in Figure 3, as Model 2 has an 80-bp trend downward in the equity premium. For Model 1 the value of the equity premium is the same at the *end* of the 53-year simulation period as it is at the *start* of the 53-year period, as Model 1 does not incorporate a downward trend or structural break in the equity premium process. Critical values of the test statistics corresponding to statistical significance at the 10%, 5%, and 1% levels are indicated by thin dotted horizontal lines in each panel, with the lowest line indicating significance at the 10% level and the highest line the 1% significance level.

Consider now specifically Graph A of Figure 3. (Note that we use a log scale for the vertical axis of the plots in Graph A of Figures 3–5 for clarity of presentation. Note as well that we postpone further discussion of Graphs B–D until after we have introduced results for Models 1–12.) On the basis of Graph A of Figure 3, only in the case of Model 4 do we observe χ^2 test statistics lower than the cutoff value implied by a 10% significance level (again, indicated by the lowest horizontal dotted line in the plot). The test statistics dip (barely) below the 10% cutoff line only for values of the equity premium within about 25 bp of 4%. Models 1–3, in contrast, are rejected at the 10% level for *every* equity premium value. If we allow fairly substantial departures of the simulated data from the S&P 500 data (e.g., test statistics that are unusual at the 1% level of significance, the upper horizontal dotted line in the plot), then all the models indicate ranges of equity premiums that are not rejected, in each case centered roughly between 3.5% and 4%. Recall that the equity premium plotted is the *ending* value, so if the model has a downward trend or decline because of a break in the equity premium, its ending value is below its average equity premium.


One conclusion to draw from the relative performance of these 4 competing models is that each additional feature over the base model (e.g., the dividend growth acceleration in the late 1970s and the trends and breaks in the equity premium) leads to better performance relative to the base model, but each is still inadequate in isolation. The model most easily rejected is clearly that which does not account for trends and breaks in the equity premium and cash flow processes.

B. Further Model Extensions (Two or More at a Time)

We turn now to joint tests based on Models 5–10. These models incorporate the basic features of Model 1, including time-varying and dependent dividend growth and interest rates, parameter uncertainty, and, with the exception of Model 10, an equity premium process derived from the Merton (1980) conditional

FIGURE 3
 Joint and Individual Test Statistics for Models 1–4

Figure 3 contains plots of test statistics for Models 1–4. Graph A plots joint χ^2 tests based on a set of 3 variables (the ex post equity premium, the mean dividend yield, and the excess return volatility) for various ending values of the ex ante equity premium for each model. In Graph A the vertical axis is plotted on a log scale. The remaining graphs contains t-Test values corresponding to tests on the individual variables for each of the models: the ex post equity premium in Graph B, the excess return volatility in Graph C, and price-dividend ratio in Graph D. In each panel the critical values of the test statistics corresponding to test significance at the 10%, 5%, and 1% levels are indicated by horizontal lines.




CAPM. These models also permit trends and/or breaks in the equity premium and dividend growth rate processes 2 or more at a time and incorporate alternative time-series models for the interest rate and the equity premium processes. Models 5–10 allow us to explore questions like: Do we need a conditional equity


FIGURE 4
 Joint and Individual Test Statistics for Models 5–8

Figure 4 contains plots of test statistics for Models 5–8. Graph A plots joint χ^2 tests based on a set of 3 variables (the ex post equity premium, the mean dividend yield, and the excess return volatility) for various ending values of the ex ante equity premium for each model. In Graph A the vertical axis is plotted on a log scale. The remaining graphs contains t-test values corresponding to tests on the individual variables for each of the models: the ex post equity premium in Graph B, the excess return volatility in Graph C, and price-dividend ratio in Graph D. In each graph the critical values of the test statistics corresponding to test significance at the 10%, 5%, and 1% levels are indicated by horizontal lines.

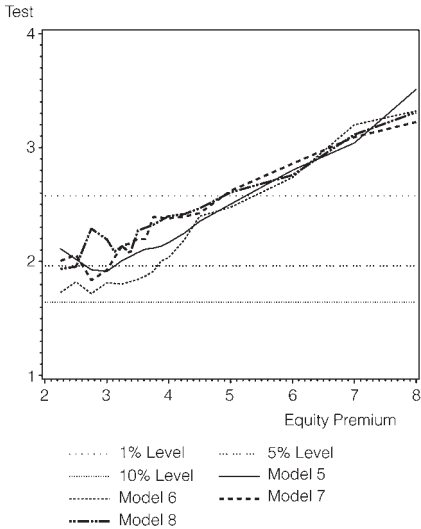
Graph A. Joint Test, All Moments



Graph B. t-Test, Ex Post Equity Premium



Graph C. t-Test, Excess Return Volatility



Graph D. t-Test, Price-Dividend Ratio


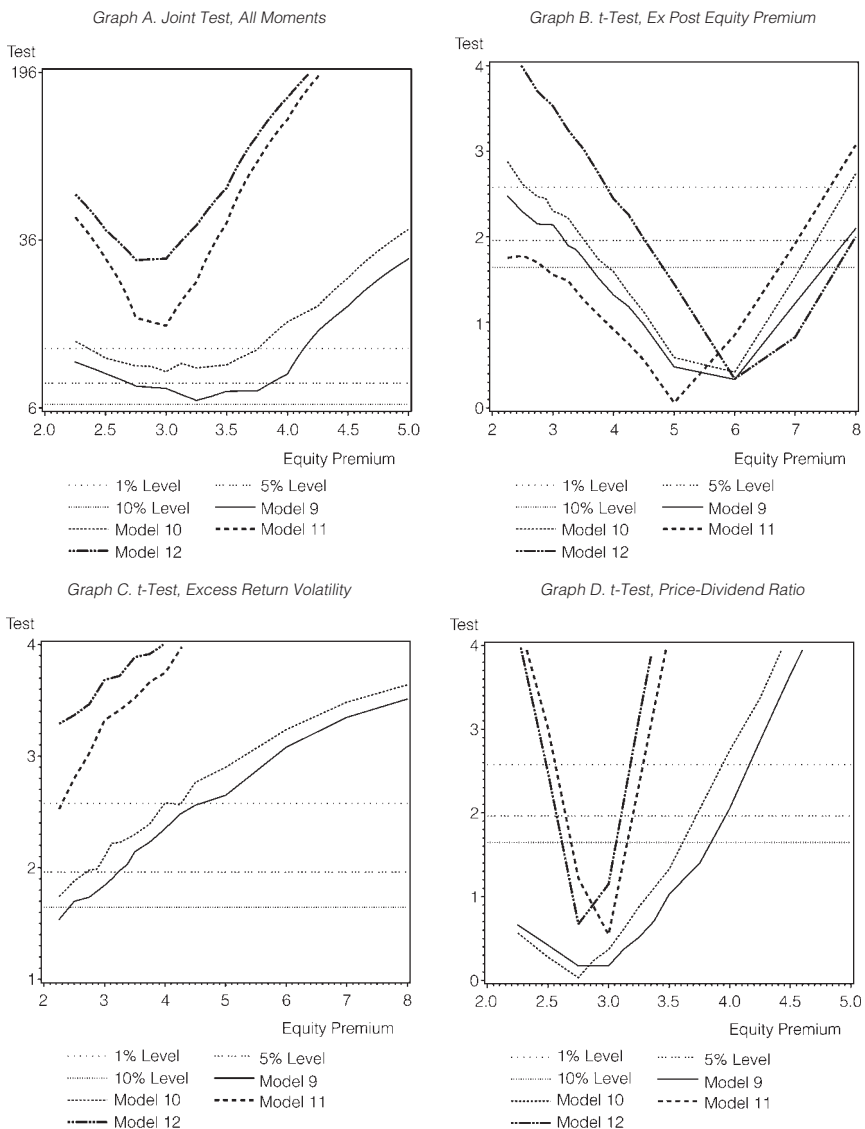


FIGURE 5
 Joint and Individual Test Statistics for Models 9–12

Figure 5 contains plots of test statistics for Models 9–12. Graph A plots joint χ^2 tests based on a set of 3 variables (the ex post equity premium, the mean dividend yield, and the excess return volatility) for various ending values of the ex ante equity premium for each model. In Graph A the vertical axis is plotted on a log scale. The remaining graphs contains t-test values corresponding to tests on the individual variables for each of the models: the ex post equity premium in Graph B, the excess return volatility in Graph C, and price-dividend ratio in Graph D. In each panel the critical values of the test statistics corresponding to test significance at the 10%, 5%, and 1% levels are indicated by horizontal lines.



premium model, or is it sufficient to have an equity premium that simply trends downward with a break? If we have a break, a trend, and time variation in the equity premium process, is it still essential to account for the disappearing dividends

of the last 25 years? Are our results sensitive to the time-series model specifications we employ?

Model 5 augments the base model (Model 1) to include an 80-bp gradual downward trend in the equity premium and a 100-bp gradual upward trend in the dividend growth rate. Model 6 is the base model adjusted to incorporate a 30-bp gradual downward trend in the equity premium, a 50-bp abrupt decline in the equity premium, and a 100-bp gradual upward trend in the dividend growth rate. Model 7 is the best model as indicated by the BIC, augmenting the equity premium process with a 30-bp gradual downward trend and a 50-bp abrupt decline and adding a 100-bp gradual upward trend in the dividend growth rate. Model 8 takes the second-best BIC model and incorporates a 30-bp gradual downward trend in the equity premium, a 50-bp abrupt decline in the equity premium, and a 100-bp gradual upward trend in the dividend growth rate. Model 9 is the base model adjusted to incorporate a 30-bp gradual downward trend in the equity premium and a 50-bp abrupt decline in the equity premium. Model 10 has the equity premium model follow a deterministic downward trend with a 50-bp structural break, interest rates follow an AR(1), and dividend growth rates follow an MA(1).

Given the existing evidence in support of a gradual downward trend in the equity premium, a structural break in the equity premium process over the early 1990s, and an increase in the growth rate of nondividend cash flows to investors (such as share repurchases) starting in the late 1970s, we believe Models 6–8 to be the best calibrated, and therefore perhaps the most plausible among all the models we consider, and Model 5 to be a close alternative.

In Graph A of Figures 4 and 5 we present plots of the χ^2 test statistics on 3 moments of the data: the mean ex post equity premium, the excess return volatility, and the mean dividend yield. Again, we consider Graphs B–D later. We see in Graph A of Figure 4 that for Models 5–8 we cannot reject a range of unconditional equity premium values at the 5% level. These models produce test statistics that drop well below even the 10% critical value. These models all embed the increased cash flow feature and either an 80-bp downward trend in the equity premium, or both a break and a trend in the equity premium adding up to an 80-bp decline over the last half century. The range of equity premiums supported (not rejected) is narrowest for Model 7 (the best model indicated by BIC) and Model 8 (the second best model indicated by BIC) with a range of less than 75 bp at the 10% level. The range is slightly wider for Models 5 and 6, roughly 75 bp–100 bp. In each case, the equity premium that yields the minimum joint test statistic, corresponding to our estimate of π , is centered between 3.25% and 3.75%.

For the models that exclude the cash flow increase, Models 9 and 10, displayed in Figure 5, we can reject all unconditional equity premium values at the 10% level. Model 9 is best compared to Model 6, as it is equivalent to Model 6 with the sole difference of excluding the cash flow increase. Graph A of Figures 4 and 5 shows that excluding the cash flow increase flattens the trough of the plot of χ^2 statistics, and approximately doubles the test statistic value, from a little over 3 for Model 6 in Figure 4 to a little over 6 for Model 9 in Figure 5. Model 10 is identical to Model 9 apart from the sole difference that Model 10 excludes

the conditionally varying equity premium process. Exclusion of this conditional variation (modeled as a 1st-order autoregressive process) worsens the ability of the model to match moments to the U.S. experience at every value of the equity premium. The difference in performance leads us to reject a model excluding a conditionally varying equity premium.

On the basis of our most plausible models, Models 6–8, we can conservatively conclude that the unconditional equity premium is within 50 bp of 3.5%. We can also conclude that models that allow for breaks and/or trends in the equity premium process are the only models that are not rejected by the data. Simple equity premium processes, those that rule out any one of a downward break and/or trend or a conditionally varying equity premium process, cannot easily account for the observed low dividend yields, high returns, and high return volatility. Ignoring the impact of share repurchases on cash flows to investors over the last 25 years also compromises our ability to match the experience of U.S. prices and returns of the last half century.

C. Is Sampling Variability (Uncertainty) in Generating Parameters Important?

All of the models we have considered so far, Models 1–10, incorporate parameter value uncertainty. This uncertainty is measured using the estimated covariance of the parameter estimates from our models. We generate model parameters by randomly drawing values from the joint distribution of the parameters, exploiting the asymptotic result that our FIML procedure produces parameter estimates that are jointly normally distributed, with an easily computed variance-covariance structure.

Now we consider 2 models that have no parameter sampling variability built into them, Models 11 and 12. In these models the point estimates from our ARMA estimation on the S&P 500 data are used for each and every simulation. Ignoring uncertainty about the true values for the parameters of the ARMA processes for interest rates, dividend growth rates, and the equity premium should dampen the variability of the generated financial statistics from these simulations, and potentially understate the range of unconditional equity premiums supported by the last half century of U.S. data. Model 11 is the base model augmented to incorporate a 30-bp gradual downward trend in the equity premium, a 50-bp abrupt decline in the equity premium, and a 100-bp gradual upward trend in the dividend growth rate, with no parameter uncertainty. (Model 11 is identical to Model 6 apart from ignoring parameter uncertainty.) Model 12 is the base model with no parameter uncertainty.

We see in Graph A of Figure 5 that both Models 11 and 12 are rejected for all values of the equity premium, though Model 11, which allows for trends and breaks, performs better than Model 12. The log scale for the vertical axis compresses the values, but the minimum χ^2 statistic for Model 12 is close to 30, indicating very strong rejection of the model, while the minimum χ^2 statistic for Model 11 is roughly 10. It is apparent that parameter uncertainty is an important model feature. Ignoring parameter uncertainty leads to model rejection, even at the equity premium setting that corresponds to the minimum test statistic.

D. The Moments That Matter

An interesting question that arises with regard to the joint tests is, where does the test power come from? That is, which variables give us the power to reject certain ranges of the unconditional equity premium in our joint χ^2 tests? An examination of the ranges of the unconditional equity premium consistent with the *individual* moments can shed some light on the source of the power of the joint tests. Graphs B, C, and D of Figures 3–5 display plots of the univariate t -test statistics based on each of the variables we consider in the joint tests plotted in Graph A of these figures. Graph B of each figure plots t -test statistics on the ex post equity premium, Graph C of each figure plots t -test statistics on the excess return volatility, and Graph D of each figure plots t -test statistics on the price-dividend ratio.

Consider first Graph B of Figures 3–5. Virtually all of the models have a minimum t -test statistic at a point that is associated with an equity premium close to 6%.¹⁸ Because our method involves minimizing the distance between the ex post equity premium based on the realized S&P 500 value (which is a little over 6%) and the ex post equity premium estimate based on the simulated data, it is not surprising that the minimum distance is achieved for models when they are set to have an equity premium close to 6%. The t -test on the mean ex post equity premium rises linearly as the equity premium setting departs from 6% for each model. The t -test does not typically reject equity premium values at the 10% level until the equity premium deviates quite far from 6%. For example, in Graph B of Figure 4 the ending equity premium must be as low as 2.25% or as high as 7% before we see a rejection at the 10% level. This wide range reflects the imprecision of the estimate of the ex post equity premium, which is also evident in the realized S&P 500 data.

The t -tests on the excess return volatility, presented in Graph C of Figures 3–5, indicate that lower equity premium values lead to models that are better able to match the S&P 500 experience of volatile returns.¹⁹ Note that as the unconditional equity premium decreases, the volatility of returns *increases*, so high equity premiums lead to simulated return volatilities that are much lower than the realized S&P 500 return volatility we have witnessed over the last half

¹⁸Recall that the equity premium values shown on the horizontal axes are *ending* values, so if the model has a downward trend or break in the equity premium process, its ending value is below the mean equity premium. For instance, Model 11 has a data generating process that incorporates trends and breaks that lead to an ending equity premium lower than the starting value. For Model 11 we observe (in Graph B of Figure 5) a minimum t -test at an *ending* value of the equity premium that is below the 6% *average* equity premium. The coarseness of the grid of equity premium values around 6% prevents this feature from being more obvious for some of the other models.

¹⁹The intuition behind this result is easiest to see by making reference to the Gordon (1962) constant dividend growth model, $P = D/(r - g)$. As the discount rate, r , declines in magnitude, the Gordon price increases. The variable r equals the risk-free rate plus the equity premium in our simulations, so low values of the equity premium lead to values of the discount rate that are closer to the dividend growth rate, resulting in higher prices. When the value of the equity premium is low, small increases in the dividend growth rate or small decreases in the risk-free rate lead to large changes in the Gordon price. In our simulations (where the conditional mean dividend growth rate and conditional mean risk-free rate change over time), when the value of the equity premium is low, small changes in the conditional means of dividend growth rates or risk-free rates also lead to large price changes (i.e., volatility).

century. The test statistic, however, rises slowly as the unconditional equity premium grows larger, in contrast to the joint test statistics plotted in Graph A of Figures 3–5, in which the χ^2 test statistic rises *sharply* as the unconditional equity premium grows larger (recall that the Graph A vertical axis has a compressed log scale in Figures 3–5). Given these contrasting patterns, the return volatility moment, *by itself*, is unlikely to cause the sharply rising joint test statistic.

Consider now the *t*-test statistics on the price-dividend ratio, plotted in Graph D of Figures 3–5. Notice that in all cases the *t*-test on the price-dividend ratio jumps up sharply as the unconditional equity premium rises above 3%. Thus the sharply increasing χ^2 statistics we saw in Graph A of the 3 figures are likely due in large part to information contained in the price-dividend ratio. However, return volatility reinforces and amplifies the sharp rejection of premiums above 4%. In terms of the 3 moments we have considered in the joint χ^2 and univariate *t*-test statistics, it is evident that the upper range of equity premiums consistent with the experience of the last half century in the U.S. is limited by the high average S&P 500 price-dividend ratio together with the high volatility of returns. This result is invariant to the way we model dividend growth, interest rates, or the equity premium process. Even an average equity premium of 5% produces economies with price-dividend ratios and return volatilities so low that they are greatly at odds with the high return volatility and high average price-dividend ratio observed over the past half century in the U.S.

E. Sensitivity to Declining Dividends Through Use of the Price-Dividend Ratio

To ensure that our results are not driven by a single moment of the data, in particular a moment of the data possibly impacted by declining dividend payments in the U.S., we perform 2 checks. First, in Models 4–8 we incorporate higher dividends and dividend growth rates than observed in U.S. corporate dividends. This is to adjust for the practice, adopted widely beginning in the late 1970s, of U.S. firms delivering cash flows to investors in ways (such as share repurchases) that are not recorded as corporate dividends. As we previously reported, Models 4–8 (the models that incorporate higher cash flows to investors than recorded by S&P 500 dividend payments, i.e., the models that use cash flows *including* share repurchases) are best able to account for the observed U.S. data. Reassuringly, the estimate of the equity premium emerging from Models 4–8 is virtually identical to that produced by the models that exclude share repurchases.

Our second check is to perform joint tests excluding the price-dividend ratio. Any sensitivity to mismeasurement of the price-dividend ratio should be mitigated if we consider joint test statistics that are based only the ex post equity premium and return volatility, excluding the price-dividend ratio. These (un-tabulated) joint tests confirm 2 facts. First, when the joint tests exclude the price-dividend ratios, the value of the χ^2 statistic rises less sharply for values of the unconditional equity premium above 4%. Essentially, this indicates that using 2 moments of the data (excluding the price-dividend ratio) rather than all 3 moments makes it more difficult to identify the minimum test statistic value and thus more difficult to identify our estimate of the equity premium. This confirms our earlier intuition that the

price-dividend ratio is instrumental in determining the steep rise of the joint test statistic in Graph A of Figures 3–5. Second, and most importantly, the minimum test statistic is still typically achieved for models with an equity premium value between 3% and 4%. For some of the models, the minimum test statistic is 25 bp or 50 bp lower than that found when basing joint tests on the full set of 3 moments. For a few models, the minimum test statistic is 25 bp or 50 bp higher. Again, Models 1–3 are rejected for every value of the equity premium, and again for Models 4–8 the range of equity premiums that are not rejected is narrow.

F. Investors' Model Uncertainty

We have been careful to explore the impact of estimation uncertainty by simulating from the sampling distribution of our model parameters, and to explore the impact of model specification choice (and implicitly model *misspecification*) by looking at a variety of models for interest rates, dividend growth rates, and equity premium rates, ranging from constant rate models to various ARMA specifications, with and without trends and breaks in the equity premium and dividend growth rates. Comparing distributions of financial statistics emerging from this range of models to the outcome observed in the U.S. over the last half century leads us to the conclusion that the range of equity premiums that could have generated the U.S. experience is fairly narrow, falling within 50 bp of 3.5%. However, we have not yet addressed the impact of *investor* uncertainty regarding the true fundamental value of the assets being priced. Up to this point, all simulated prices and returns have been generated with knowledge of the (fundamental) processes generating interest rates and dividends.

It is impossible to be definitive in resolving the impact of investor uncertainty on prices and returns. To do so we would have to know what (incorrect) model of fundamental valuation investors are actually using. We can nonetheless focus our attention on procedures likely to be less affected by investor uncertainty than others. Up to this point, the joint tests we have used to identify the plausible range of equity premiums have employed the observed return volatility over the last half century in the U.S. and the volatility of returns produced in our simulated economies. However, investor uncertainty could cause market prices to over- and under-shoot fundamental prices, impacting return volatility, perhaps significantly. A joint test statistic based on only the mean equity premium and the mean price-dividend ratio, however, should be relatively immune to the impact of investor uncertainty. (In the absence of extended price bubbles, *mean* yields should not be impacted greatly by temporary pricing errors.) Thus we consider the joint χ^2 test statistic based on only the mean return and the mean price-dividend ratio. In the interest of brevity we do not include the associated plots or statistics; however, our findings can be summarized as follows: Overall, the value of the unconditional equity premium at which the joint test statistic is minimized (i.e., our estimate of the equity premium) is not particularly affected by our having based the joint tests on these 2 moments of the data rather than the original 3 moments, nor is our selection of plausible models for the equity premium process. Across the models, the highest estimate of the unconditional equity premium is roughly 4% (for Model 4) and the lowest is 3% (for Models 11 and 12). With the joint tests based

on these 2 moments, all models support (i.e., do not reject) broader ranges of the equity premium, with the range widest for Models 4–8 (now spanning roughly 200 bp for any given model, from equity premium values as low as 2.25% for Model 7 to values as high as 4.5% for Model 4). This widening of the range of plausible equity premiums is consistent with a decline in the power of our joint test, presumably the result of omitting an important moment of the data, the return volatility. The widening of the range of plausible equity premiums is also consistent with investors being uncertain about the true fundamental value of the assets being priced. The last half century of data from the U.S. will be less informative as investor uncertainty about the processes governing fundamentals exaggerates the volatility of returns and hence reduces the precision of estimates of the equity premium.

To the extent that market prices are set in an efficient market and determined by participants with models of dividend growth rates and interest rates that reflect reality, these ranges of plausible equity premiums based on only the 2-moment joint test are overly wide. Still, these ranges are useful for putting a loose bound on the likely range of the equity premium.

V. Conclusions

The equity premium of interest in theoretical models is the extra return investors anticipate when purchasing risky stock instead of risk-free debt. Unfortunately, we do not observe this *ex ante* equity premium in the data. We only observe the returns that investors receive *ex post*, after they purchase the stock and hold it over some period of time during which random economic shocks impact prices. U.S. stocks have historically returned roughly 6% more than risk-free debt; *ex post* estimates provided by recent papers suggest the U.S. equity premium may have fallen in recent years. However, all of these estimates are imprecise, and there is little consensus emerging about the true value of the equity premium. The imprecision and lack of consensus hamper efforts to use equity premium estimates in practice, for instance to conduct valuation or to perform capital budgeting. The imprecision of equity premium estimates also complicates the resolution of the equity premium puzzle and makes it difficult to determine if the equity premium has in fact been changing over time.

In order to determine the most plausible value of the equity premium and the most plausible restrictions on how the equity premium evolves over time, we have exploited information not just on the *ex post* equity premium and the precision of this estimate, but also on related financial statistics that define the era in which this *ex post* equity premium was estimated. The idea of looking at related fundamental information to improve the estimate of the equity premium follows recent work on the equity premium that has also sought improvements through the use fundamental information like the dividend and earnings yields, higher-order moments of the excess return distribution, and return volatility and price movement directions.

Our central insight is that the knowledge that a low dividend yield, high excess return, and high return volatility all occurred together over the last 5 decades tells us something about the equity premium and the likelihood that the equity

premium is conditionally varying with trends and breaks. Certainly, if sets of these financial statistics are considered together, we should be able to estimate the equity premium more accurately than if we were to look only at the realized excess return. This insight relies on the imposition of some structure from economic models, but our result is quite robust to a wide range of model structures, lending confidence to our conclusions.

We reject as inconsistent with the U.S. experience all but a narrow range of values of the equity premium and all but a small number of equity premium models. We do so while incorporating model estimation uncertainty and allowing for investor uncertainty about the true state of the world. The range of unconditional equity premiums that is most plausible is centered very close to 3.5% for virtually every model we consider. The models of the equity premium not rejected by our model specification tests (i.e., consistent with the experience of the U.S. over the last half century) incorporate substantial autocorrelation, a structural break, and/or a gradual downward trend in the conditional equity premium process.

References

- Ang, A., and J. Liu. "A General Affine Earnings Valuation Model." *Review of Accounting Studies*, 6 (2001), 397–425.
- Ang, A., and J. Liu. "Risk, Return, and Dividends." *Journal of Financial Economics*, 85 (2007), 1–38.
- Bagwell, L. S., and J. B. Shoven. "Cash Distributions to Shareholders." *Journal of Economic Perspectives*, 3 (1989), 129–140.
- Bansal, R., and A. Yaron. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." *Journal of Finance*, 59 (2004), 1481–1509.
- Barro, R. J. "Rare Events and the Equity Premium." NBER Working Paper W11310 (2005).
- Bekaert, G., and C. R. Harvey. "Time-Varying World Market Integration." *Journal of Finance*, 50 (1995), 403–444.
- Bonaparte, Y. "The Equity Premium Puzzle: A Reconciliation." Working Paper, University of Texas (2008).
- Campbell, J. Y., and R. J. Shiller. "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors." *Review of Financial Studies*, 1 (1988), 195–228.
- Claus, J., and J. Thomas. "Equity Premia as Low as Three Percent? Evidence from Analysts' Earnings Forecasts for Domestic and International Stock Markets." *Journal of Finance*, 56 (2001), 1629–1666.
- Cochrane, J. H. *Asset Pricing*. Princeton, NJ: Princeton University Press (2001).
- Corradi, V., and N. R. Swanson. "Bootstrap Specification Tests for Diffusion Processes." *Journal of Econometrics*, 124 (2005), 117–148.
- Dimson, E.; P. Marsh; and M. Staunton. "The Worldwide Equity Premium: A Smaller Puzzle." In *Handbook of the Equity Risk Premium*, R. Mehra, ed. Amsterdam, The Netherlands: Elsevier (2008).
- Donaldson, R. G., and M. J. Kamstra. "A New Dividend Forecasting Procedure That Rejects Bubbles in Asset Prices: The Case of 1929's Stock Crash." *Review of Financial Studies*, 9 (1996), 333–383.
- Duffie, D., and K. J. Singleton. "Simulated Moments Estimation of Markov Models of Asset Prices." *Econometrica*, 61 (1993), 929–952.
- Fama, E. F., and K. R. French. "Disappearing Dividends: Changing Firm Characteristics or Lower Propensity to Pay." *Journal of Financial Economics*, 60 (2001), 3–43.
- Fama, E. F., and K. R. French. "The Equity Premium." *Journal of Finance*, 57 (2002), 637–659.
- Gordon, M. J. *The Investment, Financing and Valuation of the Corporation*. Homewood, IL: Irwin (1962).
- Graham, J. R., and C. R. Harvey. "The Long-Run Equity Risk Premium." *Finance Research Letters*, 2 (2005), 185–194.
- Hannan, E. J. "The Estimation of the Order of an ARMA Process." *Annals of Statistics*, 8 (1980), 1071–1081.
- Hodrick, R. J. "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement." *Review of Financial Studies*, 5 (1992), 357–386.

- Hull, J. C. *Options, Futures, and Other Derivative Securities*. Englewood Cliffs, NJ: Prentice-Hall (1993).
- Jorion, P., and W. N. Goetzmann. "Global Stock Markets in the Twentieth Century." *Journal of Finance*, 54 (1999) 953–980.
- Kendall, M. G., and A. Stuart. *The Advanced Theory of Statistics*, Vol. 1, 4th ed. New York, NY: MacMillan Publishing (1977).
- Kocherlakota, N. R. "The Equity Premium: It's Still a Puzzle." *Journal of Economic Literature*, 34 (1996), 42–71.
- Maheu, J. M., and T. H. McCurdy. "How Useful Are Historical Data for Forecasting the Long-Run Equity Return Distribution?" *Journal of Business and Economic Statistics*, 27 (2009), 95–112.
- McFadden, D. "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration." *Econometrica*, 57 (1989), 995–1026.
- Mehra, R. "The Equity Premium: Why Is It a Puzzle?" *Financial Analysts Journal*, 59 (2003), 54–69.
- Mehra, R., and E. C. Prescott. "The Equity Premium: A Puzzle." *Journal of Monetary Economics*, 15 (1985), 145–161.
- Mehra, R., and E. C. Prescott. "The Equity Premium in Retrospect." In *Handbook of the Economics of Finance*, Vol. 1B, G. M. Constantinides, M. Harris, and R. Stulz, eds. Amsterdam, The Netherlands: Elsevier BV (2003).
- Merton, R. C. "On Estimating the Expected Return on the Market: An Exploratory Investigation." *Journal of Financial Economics*, 8 (1980), 323–361.
- Nelson, D. B. "Conditional Heteroskedasticity in Asset Returns: A New Approach." *Econometrica*, 59 (1991), 347–370.
- Newey, W. K., and K. D. West. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55 (1987), 703–708.
- Pakes, A., and D. Pollard. "Simulation and the Asymptotics of Optimization Estimators." *Econometrica*, 57 (1989), 1027–1057.
- Pástor, L., and R. F. Stambaugh. "The Equity Premium and Structural Breaks." *Journal of Finance*, 56 (2001), 1207–1239.
- Rozeff, M. S. "Dividend Yields Are Equity Risk Premiums." *Journal of Portfolio Management*, 11 (1984), 68–75.
- Shiller, R. J. "Stock Prices and Social Dynamics." *Brookings Papers on Economic Activity*, 2 (1984), 457–498.
- Siegel, J. J., and R. H. Thaler. "Anomalies: The Equity Premium Puzzle." *Journal of Economic Perspectives*, 11 (1997), 191–200.
- Welch, I. "Views of Financial Economists on the Equity Premium and on Professional Controversies." *Journal of Business*, 73 (2000), 501–537.
- Welch, I. "The Consensus Estimate for the Equity Premium by Academic Financial Economists in December 2007." Working Paper, Brown University (2008).
- White, H. *Asymptotic Theory for Econometricians*. Orlando, FL: Academic Press (1984).