Estimating the Ex Ante Equity Premium

R. Glen Donaldson
Sauder School of Business
University of British Columbia
glen.donaldson AT sauder.ubc.ca

Mark J. Kamstra
Schulich School of Business
York University
mkamstra AT schulich.yorku.ca

Lisa A. Kramer
Rotman School of Management
University of Toronto
Lkramer AT rotman.utoronto.ca

October 2007

Keywords: equity risk premium; simulated method of moments; SMM

JEL classifications: G12, C13, C15, C22

We have benefited from the suggestions of Wayne Ferson, Ian Garrett, Mark Fisher, Joel Hasbrouck, Raymond Kan, Patrick Kelly, Alan Kraus, Tom McCurdy, Federico Nardari, Cesare Robotti, Jacob Sagi, Tan Wang, participants of the Western Finance Association Meetings, the Northern Finance Association Meetings, the Canadian Econometrics Study Group, the European Econometric Society Meetings, the investment conference of the University of Colorado at Boulder's Burridge 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 US Securities & 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. A previous version of this paper was titled "Stare Down the Barrel and Center the Crosshairs: Targeting the Ex Ante Equity Premium." The most recent version of this paper is available from: http://ssrn.com/abstract=945192.

Estimating the Ex Ante Equity Premium

Abstract

We find that the true ex ante equity premium very likely lies within 50 basis points of 3.5%. This estimate is similar to values obtained in some recent studies but is considerably more precise. In addition to narrowing the range of plausible ex ante equity premia, 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 US 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. Our approach involves simulating the distribution from which interest rates, dividend growth rates, and equity premia are drawn and determining the prices and returns consistent with these distributions. We achieve the narrower range of ex ante equity premium values and the narrower set of plausible models by comparing statistics that arise from our simulations with key financial characteristics of the US economy, including the mean dividend yield, return volatility, and mean return. Our findings are achieved in part with the imposition of more structure than is typically exploited in the literature. In order to mitigate the potential for misspecification with this additional structure, we consider a broad collection of models that variously do or do not incorporate features such as an adjustment in dividend growth rates to account for recently increased share repurchase activity, sampling uncertainty in generating model parameters, and cross-correlation between interest rates, dividend growth rates, and equity premia.

Estimating the Ex Ante Equity Premium

Financial economic theory is often concerned with the premium that investors demand ex ante, when they first decide whether to purchase risky stocks instead of risk-free debt. In contrast, empirical tests of the equity premium often focus on the return investors received ex post.[1] It is well known that estimates of the ex ante equity premium based on ex post data can be very imprecise; such estimates have very wide margins of error, as wide as 1000 basis points in typical studies and 320 basis points in some recent studies. This fact makes it challenging to employ the equity premium estimates for common practical purposes, including evaluating the equity premium puzzle, performing valuation, and conducting capital budgeting. The imprecision of traditional equity premium estimates also makes it difficult to determine if the equity premium has changed over time. Our goals, therefore, are to develop a more precise estimate of the ex ante equity premium and to determine what kind of equity premium model can be supported by the experience of US markets. We accomplish these goals by employing simulation techniques that identify a range of models of the equity premium and the values of the ex ante equity premium that are consistent with values of several key financial statistics that are observed in US market data, including dividend growth rates, interest rates, Sharpe ratios, price-dividend ratios, volatilities, and of course the ex post equity premium.

Our results suggest that the mean ex ante equity premium lies within 50 basis points of 3.5%. These results stand even when we allow for investors' uncertainty about the true state of the world. The tightened bounds are achieved in part with the imposition of more structure than has been commonly employed in the equity premium literature. In order to mitigate the potential for misspecification with this additional structure, we consider a broad collection of models that variously do or do not incorporate features such as a conditionally time-varying equity premium, a downward trend in the equity premium, a structural break in the equity premium, an adjustment in dividend growth rates to account for increased share repurchase activity in the last 25 years, sampling uncertainty in generating model parameters, a range of time series models, and cross-correlation between interest rates, dividend growth rates, and equity premia. We also find that

[1] The 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), and Jain (2005). For excellent surveys see Kocherlakota (1996), Siegel and Thaler (1997), Mehra and Prescott (2003), and Mehra (2003).

equity premium models that allow for time-variation, breaks, and/or trends in the equity premium process are the models that best match the experience of US 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, itself an important finding.

We draw on two relatively new techniques in order to provide a more precise estimate of the equity premium than is currently available. The first technique builds on the fundamental valuation dividend discounting method of Donaldson and Kamstra (1996). This technique permits the simulation of fundamental prices, returns, and return volatility for a given ex ante equity premium. Donaldson and Kamstra find that if we allow dividend growth rates and discount rates 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. The second technique is simulated method of moments (SMM).[2] An attractive feature of SMM is that the estimation of parameters requires only that the model, with a given set of parameters, can generate data. SMM forms estimates of model parameters by using a given model with a given set of parameter values to simulate moments of the data (for instance means or volatilities), measuring the distance between the simulated moments and the actual data moments, and repeating with new parameter values until the parameter values that minimize the (weighted) distance are found.[3] 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, Section 11.6). We use SMM rather than GMM because, as we show below, the economic model we use is nonlinear in the parameters and cannot be solved without the use of SMM.

We exploit the dividend discounting method of Donaldson and Kamstra to generate simulated fundamental prices, dividends, returns, and derivative moments such as the mean ex post equity

[2]Simulated method of moments was developed by McFadden (1989) and Pakes and Pollard (1989), and a helpful introduction to the technique is provided in Carrasco and Florens (2002). Examples of papers that employ SMM in an asset pricing context are Duffie and Singleton (1993) and Corradi and Swanson (2005).

[3]The typical implementation of SMM is to weight the moments inversely to their estimated precision; that is minimize the product of the moments weighted by the inverse of the covariance matrix of the moments. This is the approach we adopt.

premium, mean dividend yield, and return volatility for a given ex ante equity premium. We minimize (by choice of the ex ante equity premium) the distance between the simulated moments that the model produces and the moments observed in US stock markets over the past half century. That is, given various characteristics of the US economic experience (such as low interest rates and a high ex post equity premium, high Sharpe ratios and low dividend yields, *etc.*), we determine the range of values of the ex ante equity premium and the set of equity premium models that are most likely to have generated the observed collection of sample moments.

To undertake our study, we consider a broad collection of models, including models with and without conditional time-variation in the equity premium process, with and without trends in the equity premium, with and without breaks in the equity premium, with and without breaks in the dividend growth rate, as well as various autoregressive specifications for dividend growth rates, interest rates, and the equity premium. Virtually every model we consider achieves a minimum distance between the simulated moments and the actual data moments by setting the ex ante equity premium between 3% and 4%, typically very close to 3.5%. That is, the equity premium estimate is very close to 3.5% across our models. Further, the range of ex ante equity premium values that can be supported by the US data for a given model is typically within plus or minus 50 basis points of 3.5%. Our models of fundamentals, which capture the dynamics of actual US dividend and interest rate data, imply that the true ex ante equity premium is 3.5% plus or minus 50 basis points. Simpler models of fundamental valuation, such as the Gordon (1962) constant dividend growth model, are overwhelmingly rejected by the data. Models of the equity premium which do not allow time-variation, trends, or breaks are also rejected by the SMM model specification tests. While we restrict our attention to a stock market index in this study, the technique we employ is more broadly applicable to estimating the equity premium of an individual firm.

In the literature to date, empirical work investigating the equity premium has largely consisted of a series of innovations around a common theme: producing a better estimate of the mean ex ante equity premium. Recent work in the area has included insights such as exploiting dividend yields or earnings yields to provide new, more precise estimates of the return to holding stocks (see Fama and French, 2002, and Jagannathan, McGrattan, and Scherbina, 2000), looking across many countries to account for survivorship issues (see 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 (see Dimson, Marsh, and Staunton, 2007), modeling equity premium structural breaks in a Bayesian econometric framework (see Pástor and Stambaugh, 2001), or computing out-of-sample forecasts of the distribution of excess returns, allowing for structural breaks which are identified in real time (see Maheu and McCurdy, 2007). Most of this work estimates the ex ante equity premium by considering one moment of the data at a time, typically the mean difference between an estimate of the return to holding equity and a risk free rate, though Maheu and McCurdy (2007) consider higher-order moments of the excess return distribution and Pástor and Stambaugh (2001) incorporate return volatility and direction of price movements through their use of priors.

Unfortunately, the equity premium is still estimated without much precision. Pástor and Stambaugh (2001), exploiting extra information from return volatility and prices, narrow a two standard deviation confidence interval around the value of the ex ante equity premium to plus or minus roughly 280 basis points around a mean premium estimate of roughly 4.8% (a range that spans 2% to 7.6%) and determine that the data strongly support at least one break in the equity premium in the last half century. Fama and French (2002), based on data from 1951 to 2000, provide point estimates of the ex post equity premium of 4.32% (based on earnings growth rate fundamentals) plus or minus roughly 400 basis points (again, two standard deviations) and of 2.55% (based on dividend growth rate fundamentals) plus or minus roughly 160 basis points: a range of approximately 0.95% to 4.15%. That is, the plausible range of equity premia that emerge from Fama and French's study occupy a confidence bound with a width of anywhere from 320 to 800 basis points. Claus and Thomas (2001), like Fama and French (2002), make use of fundamental information to form lower estimates of the ex post equity premium, but their study covers a shorter time period relative to the Fama and French study – 14 years versus 50 years – yielding point estimates that are subject to at least as much variability as the Fama and French estimates.

Not only are the point estimates from the existing literature imprecisely estimated in terms of their standard error, there is also less of an emerging consensus than one would hope. Fama and French (2002) produce point estimates of 2.55% (using dividend yields) and 4.78% (using earnings yields), Pástor and Stambaugh (2001) estimate the equity premium at the end of the 1990s to be 4.8%, and 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). Based on a survey of US CFOs, Graham and Harvey (2005) estimate the ten-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 above all fall within two standard errors of the sample mean estimate of the equity premium, based on US data. Further, the studies that provide these estimates do not explicitly consider which models of the equity premium process can be rejected by actual data, though Pástor and Stambaugh's analysis strongly supports a model that incorporates breaks in the equity premium process.

The remainder of our paper proceeds as follows. The basic methodology of our simulation approach to estimating equity premia is presented in Section 1, along with important details on estimating the equity premium. (Appendices to the paper provide detailed explanations of the technical aspects of our simulations, including calibration of key model parameters.) In Section 2 we compare univariate financial statistics that arise in our simulations with US market data, including dividend yields, Sharpe ratios, and conditional moments including ARCH coefficients. Our results confirm that the simulations generate data broadly consistent with the US market data and, taken one-at-a-time, these financial statistics imply that the ex ante equity premium lies in a range much narrower than between 2% and 8%. We determine how much narrower in Section 3 by exploiting the full power of the simulation methodology. We compare joint multivariate distributions of our simulated data with observed US data, yielding a very precise estimate of the ex ante equity premium and providing strong rejections of models of the equity premium process that fail to incorporate time variation, breaks, and/or trends. We find the range of ex ante equity premium values is very narrow: 3.5% plus or minus 50 basis points. Our consideration of a broad collection of possible data generating processes and models lends confidence to the findings. Section 4 concludes.

I Methodology

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

$$R_t = \frac{D_{t+1} + P_{t+1} - P_t}{P_t}.$$

The risk-free rate, set at the beginning of each period, is denoted $r_{t,f}$. The ex ante equity premium, π, is defined as the difference between the expected return on risky assets, $E\{R_t\}$, and the expected risk-free rate, $E\{r_{t,f}\}$:[4]

$$\pi \equiv E\{R_t\} - E\{r_{t,f}\}. \tag{1}$$

We do not observe this ex ante equity premium. Empirically, we only observe the returns that investors actually 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 ex post equity premium is typically estimated using historical equity returns and risk-free rates. Define $\overline{R}$ as the average historical annual return on the S&P 500 and $\overline{r}_f$ as the average historical return on US T-bills. Then we can calculate the estimated ex post equity premium, $\hat{\pi}$, as follows:

$$\hat{\pi} \equiv \overline{R} - \overline{r}_f. \tag{2}$$

Given that the world almost never unfolds exactly as one expects, there is no reason to believe that the stock return we estimate ex post is exactly the same as the return investors anticipated ex ante. It is therefore difficult to argue that just because we observe a 6% ex post equity premium in the US data, the premium that investors demand ex ante is also 6% and thus a puzzling challenge to economic theory. So we ask the following question: If investors' true ex ante premium is π, what is the probability that the US 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

[4] See, for instance, Mehra and Prescott (1985), Equation (14). We will consider time-varying equity premium models below.

observed in the US data is consistent with various ex ante premium values, π, with which standard economic theory may be more compatible. We also ask a deeper question: If investors' true ex ante premium is π, what is the probability that we would observe the various *combinations* of key financial statistics and yields that have been realized in the US, 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 premia consistent with the US market data, especially relative to previous studies that have considered univariate distributions.

Because the *empirical* joint distribution of the financial statistics we wish to consider is difficult or impossible to estimate accurately, in particular the joint distribution *conditional* on various ex ante equity premium values, we use simulation techniques to estimate this distribution. The simulated joint distribution allows us to conduct formal statistical tests that a given ex ante equity premium could have produced the US experience. Most of our models employ a time-varying ex ante equity premium, so that a simulation described as having an ex ante equity premium of 2.75% actually has a mean ex ante equity premium of 2.75%, while period-by-period the ex ante equity premium can vary somewhat from this mean value. In what follows we refer to the ex ante equity premium and the mean ex ante equity premium interchangeably.

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 price of the stock is then given by Equation (3),

$$P_t = E_t\left\{\frac{D_{t+1} + P_{t+1}}{1+r_t}\right\}, \tag{3}$$

where E_t is the conditional expectations operator incorporating information available to the market when P_t is formed, up to but not including the beginning of period t (*i.e.*, information from the end of period $t-1$ and earlier).

Assuming the usual transversality conditions, we can derive Equation (4) by recursively substituting out for future prices in Equation (3):

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

Defining the growth rate of dividends over the period t as $g_t \equiv (D_{t+1} - D_t)/D_t$, we can re-write Equation (4) as

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

Hence we can re-write Equation (1) as

$$\pi \equiv E \left\{ \frac{D_{t+1} + D_{t+1} E_{t+1} \left\{ \sum_{j=0}^{\infty} \Pi_{i=0}^{j} \frac{1+g_{t+1+i}}{1+r_{t+1+i}} \right\} - D_t E_t \left\{ \sum_{j=0}^{\infty} \Pi_{i=0}^{j} \frac{1+g_{t+i}}{1+r_{t+i}} \right\}}{D_t E_t \left\{ \sum_{j=0}^{\infty} \Pi_{i=0}^{j} \frac{1+g_{t+i}}{1+r_{t+i}} \right\}} - r_{t,f} \right\} \tag{6}$$

or

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

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

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

Under interesting conditions, such as risk-free rates and dividend growth rates that conditionally time-vary and covary (we consider, for instance, ARMA models and correlated errors for dividend growth rates and interest rates), the individual conditional expectations in Equation (8) are analytically intractable. The difference between the sample mean return and the sample mean risk-free interest rate provides a consistent estimate of π, as shown by Mehra and Prescott (1985), but unfortunately the sample mean difference is very imprecisely estimated, even based on more than 100 years of data.

We note that another consistent estimator of π is one that directly exploits the method of Donaldson and Kamstra (1996), hereafter referred to as the DK method. The DK method uses (ARMA) models for dividend growth rates and interest rates to simulate the conditional expectations $E_t\left\{\sum_{j=0}^{\infty}\Pi_{i=0}^{j}\frac{1+g_{t+i}}{1+\pi+r_{t+i,f}}\right\}$ and $E_{t+1}\left\{\sum_{j=0}^{\infty}\Pi_{i=0}^{j}\frac{1+g_{t+1+i}}{1+\pi+r_{t+1+i,f}}\right\}$. The DK method allows us, for a given *ex ante* equity premium (or time-varying equity premium process), to simulate the conditional expectations in Equation (8) as well as related (unconditional) moments, including the expected dividend yield, return volatility, ex post equity premium, and Sharpe ratio. 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 procedure) and the analogous sample moments (from the US experience over the last half century). The estimation of these expectations relies on the exact form of the conditional models for dividend growth rates and interest rates, that is, the parameters that characterize these models. A joint estimation of these models' parameters and π (*i.e.* minimizing the distance between simulated and sample moments by varying all the model's parameters and π at once) would be computationally very difficult. We utilize a two-step procedure in which first, for a given ex ante 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 ex ante equity premium to find our SMM estimate of π.

It is helpful to consider some examples of estimators based on our simulation technique. The simplest estimator would have us considering only the ex ante equity premium moment, $\pi = E\left[R_t\right] - E\left[r_{f,t}\right]$, ignoring other potentially informative moments of the data, such as the dividend yield and return volatility. Exploiting the DK procedure, we would find that the π in Equation (8) which matches the ex post equity premium (the sample moment analogue of Equation (8)) is the sample estimate of the ex post equity premium, roughly 6%. That is, in this simplest case, when we minimize the distance between the sample moment and the simulated moment and find that the estimate of the ex ante equity premium is the ex post equity premium, we do so by construction. If the DK method is internally consistent, and if we are fitting only the ex post equity premium sample moment, then the difference must be zero at the value of π equal to the ex post equity premium. This DK estimator of π, considering only one moment of the data, would offer no advantage over the ex post equity premium, which is the traditional estimate of the ex ante equity premium. Adding a second moment to our estimation procedure, say the dividend yield, and minimizing the distance between the simulated and sample moments for the ex post equity premium and the dividend yield *jointly*, would likely lead to a somewhat different ex ante equity premium estimate. Furthermore, the estimate would be more precisely estimated (*i.e.,* with a smaller standard error) since two moments are exploited to estimate the ex ante equity premium, not just one moment, at least if the extra moment of the data provided some unique information about the value of the parameter π.

The DK method provides simulated dividend yields, ex post equity premia, and any other statistic that is derivative to returns and prices, such as return volatility, resulting in a broad collection of simulated moments with which to compare moments of the actual US data in order to derive an estimator. The large collection of available moments makes it likely that our analysis can provide a tighter bound on the value of the ex ante equity premium than has been achieved previously.

B The Simulation

To estimate the joint distribution of the financial quantities of interest, we consider models calibrated to the US economy. (We calibrate to US data over 1952 through 2004, with the starting year of 1952 motivated by the US Federal Reserve Board's adoption of a modern monetary policy regime in 1951.) We provide specific details on the nature of the models we consider and how we conduct our simulations in Appendices 1 and 2. Our entire procedure can be generally summarized in the following five steps:

Step 1: Specify assumptions about the ex ante equity premium demanded by investors. Is the premium constant or time-varying? If constant, what value does it take? If time-varying, how does the value change over time? Are there any structural breaks in the equity premium process over time? 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 0.80% in total since the early 1950s. Pástor and Stambaugh (2001) also find fairly strong support for there having been a structural break over the 1990s which led to a 0.5% drop in the equity premium.[5]

Once the process driving the ex ante equity premium is defined, we can specify the discount rate (which equals the risk-free rate plus the equity premium) that an investor would rationally apply to a forecasted dividend stream in order to calculate the present value of a dividend-paying stock. Note that if the equity premium varies over time, then the models generated in the next step are calibrated to mimic the degree of covariation between interest rates, dividend growth rates, and equity premia observed in the US data.

Step 2: Estimate econometric models for the time-series processes driving actual dividends and interest rates in the US economy, allowing for autocorrelation and covariation as observed in the US data. These models will later be used to Monte-Carlo simulate a variety of potential paths for US dividends and interest rates. The simulated dividend and interest rate paths are of course different in each of these simulated economies because different sequences of random innovations are applied to the common stochastic processes in each case. However, the key drivers of the simulated economies themselves are all still identical to those of the US economy since all economies share common stochastic processes fitted to US data.

Some of the models we consider assume that all cashflows received by investors come in the form of dividends (the standard assumption). Another set of models we consider embed higher cashflows and cashflow growth rates than observed in the US S&P 500 dividend data, to account for the observation of Bagwell and Shoven (1989), Fama and French (2002), and others, that dividends under-report total cashflows to shareholders. As reported by these authors, firms have been increasingly distributing cash to shareholders via share repurchases instead of via dividends, a phenomenon commonly known as disappearing dividends, a practice adopted widely beginning in the late 1970s. Fama and French find evidence that the disappearance of dividends is in part due to an increase in the inflow of new listing to US stock exchanges, representing mostly young companies

[5]A falling equity premium is thought to come from several sources, including the declining cost of diversifying through mutual funds over the last half century, the infeasibility before the advent of mutual funds to hold fully diversified portfolios (hence higher returns required by investors to hold relatively undiversified positions), and the broader pool of investors now participating in equity ownership, sharing in the market risk and presumably lowering the required rate of return to risky assets. See Siegel (1999) and Diamond (2000).

with the characteristics of firms that would not be expected to pay dividends, and in part due to a decline in the propensity of firms to pay dividends.

Thus, for some models in our simulations, we adopt higher cashflows than would be indicated by considering US dividend data alone. On a broad set of data, Grullon and Michaely (2002) find that total payouts to shareholders have remained fairly flat, not growing over the period we consider. To the extent that this is true of the S&P 500 data, the models we consider with upward-trending dividend growth are overly aggressive, but as we show below, the higher dividend growth rate only widens the range of plausible ex ante equity premia, meaning our estimate of the precision of our approach is conservative.

Step 3: Allow for the possibility of estimation error in the parameter values for the dividend growth rate, interest rate, and equity premium time-series models. That is, incorporate into the simulations uncertainty about the true parameter values. This allows for some models with more autocorrelation in the dividend growth, interest rate, and equity premium series, some with less, some with more correlation between the processes, some with less, some with a higher variance or mean of dividend growth and interest rates, some with less, and so on. This uncertainty is measured using the estimated covariance of the parameter estimates from our models generated in Steps 1 and 2, and the procedure to randomly select parameters from the estimated joint distribution of the parameters is detailed in Appendix 1. We also account for investor uncertainty about the true fundamental processes underlying prices and returns by performing tests insensitive to this uncertainty and its impact on prices and returns, as we describe below.

Further details about Steps 1 through 3 are contained in Appendix 1. Before continuing with summarizing Steps 4 and 5 of our methodology, it is worth identifying some models that emerge from various combinations of the assumptions embedded in Steps 1 through 3. The key models we consider in this paper are shown in Table I. The first column of Table I indicates numbering that we assign to the models. The second column specifies the time-series process used to generate the interest rate and dividend growth rate series, corresponding to Step 2. The next three columns relate to Step 1 above, indicating whether or not the ex ante equity premium process incorporates a downward trend over time (and if so, how much the mean ex ante equity premium in 1952 differs from the value in 2004), whether or not there is a structural break (consisting of a 50 basis point drop) in the equity premium consistent with the findings of Pástor and Stambaugh (2001), and whether or not there is a break in the dividend growth rate process, consistent with the Bagwell and Shoven (1989) and Fama and French (2002) finding of an increase in share repurchases from the late 1970s onward.[6] The last column corresponds to Step 3, showing which models incorporate uncertainty in generating parameters. We consider a selection of 12 representative models, ranging from a simple model with no breaks or trends in the equity premium process (Model 1) to very complex models.[7] Each model is fully explored in the sections that follow. We now continue describing the two final steps of our basic methodology.

Table I goes about here.

Step 4: Calculate the fundamental stock returns (and hence ex post equity premia) that arise in each simulated economy, using a discounted-dividend-growth-rate model and based on assumptions about the ex ante equity premium from Step 1, the dividend growth rate and interest rate processes specified in Step 2, and the possible parameter uncertainty specified in Step 3. The model is rolled out to produce 53 annual observations of returns, prices, dividends, interest rates, and so on, mimicking the 53 years of annual US data available to us for comparison. Keep in mind the fact that the assumptions made in Steps 1 through 3 are the same for all simulated economies in a given experiment. That is, all economies in a given experiment have the same ex ante equity premium model (for instance a constant ex ante equity premium, or perhaps an ex ante equity premium that time-varies between a starting and ending value) and yet all economies in the set of simulations have different ex post equity premia. Given the returns and ex post equity premia for each economy, as well as the means of the interest rates and dividend growth rates produced for each economy, we are able to calculate various other important characteristics, including return volatility,

[6]In each case where we consider model specifications intended to capture real-world features like breaks and trends in rates and premia, we adopt parameterizations that bias our results to be more conservative (*i.e.* to produce a wider confidence interval for the ex ante equity premium). This allows us to avoid over-stating the gains in precision possible with our technique. For example, while Pástor and Stambaugh (2001) find evidence that there was a break in the equity premium process across several years in the 1990s, we concentrate the entire break into one year (1990). Allowing the break to be spread across several years would lead to a narrower bound on the ex ante equity premium than we find. See Appendix 1 for more details.

[7]For the sake of brevity, the Gordon (1962) constant dividend growth model is excluded from the set of models we explore in this paper. We did analyze the Gordon model and found it to perform very poorly. The model itself is rejected at every value of the ex ante equity premium, even more strongly than any other simple model considered in this paper is rejected.

dividend yields, and Sharpe ratios. There is nothing in our experimental design to exclude (rational) market crashes and dramatic price reversals. Indeed our simulations do produce such movements on occasion. The details of Step 4 are provided in Appendix 2.

Step 5: Examine the distributions of variables of interest, including ex post equity premia, Sharpe ratios, dividend yields, and regression coefficients (from estimating AR(1) and ARCH models for returns) that arise conditional on various mean values and various time-series characteristics of the ex ante equity premia. Comparing the performance of the US economy with various univariate and multivariate distributions of these quantities and conducting joint hypothesis tests allows us to determine a narrow range of equity premia consistent with the US market data. That is, only a small range of mean ex ante equity premia and time-varying equity premium models could have yielded the outcome of the past half century of high mean return and return standard deviation, low dividend yield, high ex post equity premium, *etc.*

A large literature makes use of similar techniques in many asset pricing applications, directly or indirectly simulating stock prices and dividends under various assumptions to investigate price and dividend behavior.[8] However, these studies typically employ restrictions on the dividend and discount rate processes in order to obtain prices from some variant of the Gordon (1962) model and/or some log-linear approximating framework. For instance, the present value (price, defined as P_0) of an infinite stream of expected discounted future dividends can be simplified under the Gordon model as

$$P_0 = D_1/(r - g), \tag{9}$$

where D_1 is the coming dividend, r is the constant discount rate, and g is the constant dividend growth rate. That is, by assuming constant r and g, one can analytically solve for the price. If, however, discount rates or dividend growth rates are in fact conditionally time-varying, then the infinite stream of expected discounted future dividends in Equation (5) cannot be simplified into Equation (9), and it is difficult or impossible to solve prices analytically without imposing other simplifying assumptions.

[8] See, for example, Scott (1985), Kleidon (1986), West (1988a,b), Campbell (1991), Gregory and Smith (1991), Mankiw, Romer, and Shapiro (1991), Hodrick (1992), Timmermann (1993, 1995), and Campbell and Shiller (1998).

Rather than employ approximations to solve our price calculations analytically, we instead simulate the dividend growth and discount rate processes directly, and evaluate the expectation through Monte Carlo integration techniques, adopting the DK method.[9] In the setting of time-varying dividend growth rates and interest rates which conditionally covary, this technique allows us to evaluate prices, returns, and other financial quantities without approximation error.[10] We also take extra care to calibrate our models to the time-series properties of actual market data. For example, annual dividend growth is strongly autocorrelated in the S&P 500 stock market data, counter to the assumption of a logarithmic random walk for dividends sometimes employed for tractability in other applications. Furthermore, interest rates are autocorrelated and cross-correlated with dividend growth rates. Thus we incorporate these properties in our 12 models (shown in Table I), which we use to produce our simulated dividend growth rates, interest rates, and, ultimately, our estimate of the ex ante equity premium.

We estimated each of the 12 models over a grid of discrete values of the ex ante equity premium, with the grid as fine as an eighth of a percent in the vicinity of a 3.5% equity premium, and no coarser than 100 basis points for equity premium values exceeding 5%. The entire exercise was conducted using distributed computing across a grid of 30 high-end, modern-generation computers over the course of a month. On a modern stand-alone computer, estimation of a single model for a single assumed value of the ex ante equity premium would take roughly one week to estimate (and, as stated above, we consider many values of the ex ante equity premium for each of our models).

II 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 I. Model 1 incorporates interest rates that follow an AR(1) process and dividend growth rates that follow a MA(1) process. The ex ante equity premium in Model 1 follows an AR(1) process (that emerges from Merton's (1980) conditional CAPM, as detailed in Appendix 1), with no trends or breaks in either the equity premium process or dividend growth rate process. We start with this "plain

[9]The Dondaldson and Kamstra (1996) method nests other fundamental dividend-discounting valuation methods as special cases. For instance, in a Gordon (1962) world of constant dividend growth rates and interest rates, the DK method produces the Gordon model price, albeit through numerical integration rather than analytically.

[10]There is still Monte Carlo simulation error, but that is random, unlike most types of approximation error, and it can also be measured explicitly and controlled to be very small, which we do, as explained in Appendix 2.

vanilla" model because it provides a good illustration of how well dividend-discounting models that incorporate time-varying autocorrelated dividend growth and discount rate processes can produce prices and returns that fit the experience of the last half century in the US. 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 more complex and arguably more realistic models incorporating trends and breaks later in the paper.

It is well known that the ex ante equity premium is estimated with error. See, for instance, Merton (1980), Gregory and Smith (1991), and Fama and French (1997). Any particular realization of the equity premium is drawn from a distribution, implying that given key information about the distribution (such as its mean and standard deviation), one can construct a confidence interval of statistically similar values and determine whether a particular estimate is outside the confidence interval. As mentioned above, an implication of this estimation error is that most studies have produced imprecise estimates of the mean equity premium. For instance, a typical study might yield an 800 basis point 95% confidence interval around the ex ante equity premium.[11] Studies including Fama and French (2002) have introduced innovations that make it possible to narrow the range. One of our goals is to further sharpen the estimate of the mean ex ante equity premium.

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 mean ex ante equity premium. While the analysis in this section based on univariate empirical distributions is somewhat casual, in Section III 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 mean ex ante equity premium and identifying plausible models of the equity premium process, representing our main contributions.

Consider the following: conditional on a particular value of the ex ante equity premium, how unusual is an observed realization of the ex post equity premium? How unusual is an observed realization of the mean dividend yield? Each simulated economy produces a set of financial statistics based on the simulated *annual* time-series observations, and these financial statistics can be

[11] This particular range is based on the simple difference between mean realized equity returns and the average riskfree rate based on the last 130 years of data, as summarized in Table I of Fama and French (2002).

compared and contrasted with the US experience of the last half century. By considering not only the mean of a financial statistic across simulated economies, such as the mean ex post equity premium, but also conditional moments and higher moments including the standard deviation of excess returns produced in our simulations, we can determine with high refinement the ability of our simulated data to match characteristics of the US economy. For instance, market returns, to be discussed below, are volatile. Thus it is interesting to examine the degree to which our simulations are able to produce volatile returns and to look at the distribution of return variance as we vary the mean ex ante equity premium in our simulated economies.

We can compare any financial statistic from the last half century to our simulated economies provided the statistic is based on returns or 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 two 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 ex ante equity premia. For instance, 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; moments must be finite, for instance. 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).

Rather than presenting copious volumes of tabled results, we summarize the simulation results with concise plots of probability distributions of the simulated data for various interesting financial statistics. This permits us to determine if a particular ex ante equity premium produces financial statistics similar to what has been seen over the last half century in the US.

Figure 1 contains four panels, and in each panel we present the probability distribution function for one of various financial statistics (ex post equity premia, dividend yield, Sharpe ratio, and return volatility) based on each of four different ex ante equity premium settings. We also indicate the realized value for the actual US data. Comparison of the simulated distribution with realized values in these plots permits a very quick, if casual, first assessment of how well the realized US data agree with the simulated data, and which assumed values of the ex ante equity premium appear inconsistent with the experience of the last half century of US data.

Panels A through D of Figure 1 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 mean ex ante equity premia of 2.75%, 3.75%, 5%, and 8%. For the sake of clarity, the dotted lines depicting the PDFs in Figure 1 are thinnest for the 2.75% case and become progressively thicker for the 3.75%, 5%, and 8% cases. The actual US realized data is denoted in each panel with a solid vertical line.

The actual US mean equity premium, displayed in Panel A, is furthest in the right tail of the distribution corresponding to a 2.75% ex ante equity premium, and furthest in the left tail for the ex ante premium of 8%. The wide range of the distribution of the mean ex post equity premia for each assumed value of the ex ante equity premium is consistent with the experience of the last half century in the US, in which the mean ex post equity premium has a 95% confidence interval spanning plus or minus roughly 4% or 5%. The actual dividend yield of 3.4%, displayed in Panel B, is unusually low for the 5% and 8% ex ante equity premium cases, but it is near the center of the distribution for the ex ante premium values of 2.75% and 3.75%. In Panel C, only the Sharpe ratios generated with an ex ante equity premium of 8% appear inconsistent with the US experience of the last half century. The return volatility, displayed in Panel D, clearly indicates that the experience of the US over the last half century is somewhat unusual for all ex ante equity premia considered, though least unusual for the lowest ex ante equity premium. Casual observation, based on only the evidence in these univariate plots, implies that the ex ante equity premium which could have generated the actual high ex post equity premium and low dividend yield of the last half century of the US experience likely lies above 2.75% and below 5%.

Figure 1 goes about here.

We constructed similar plots for the mean return and for conditional moments, including the return first order autocorrelation coefficient estimate (the OLS parameter estimate from regressing returns on lagged returns and a constant, *i.e.,* the AR(1) coefficient), the return first order autoregressive conditional heteroskedasticity coefficient estimate (the OLS parameter estimate from regressing squared residuals on lagged squared residuals and a constant, *i.e.,* the ARCH(1) coefficient), and the price-dividend ratio's first order autocorrelation coefficient estimate (the OLS parameter estimate from regressing the price-dividend ratio on the lagged price-dividend ratio and a constant). The mean return distributions are similar to the ex post equity premium distributions shown in Figure 1, and all choices of the ex ante equity premium produce returns and price-dividend ratios that have conditional time-series properties matching the US data, so these results are not presented here.

Figure 1 has two 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 US economy over the past five decades. Most simulated statistics match the magnitudes of financial quantities from the actual US data, even though we do not calibrate to prices or returns.[12] Second, the results suggest that the 2.75% through 8% interval we present here likely contains the ex ante equity premium consistent with the US economy. Univariate results for Models 2 through 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 US economy. Detailed univariate results for Models 2 through 12 are omitted for the sake of brevity, but the poor performance of Models 11 and 12 will be evident in multivariate results reported below.

To narrow further the range of plausible ex ante equity premium values, we need to 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 observed data. We consider the multivariate distributions of several moments of the data, including ex post equity premia, dividend yields, and return volatility. This exercise allows for inference that is not feasible with the univariate analysis conducted above, and it leads to a very precise estimate of the ex ante equity premium. We turn to this task in the next section, where we also broaden the class of models we consider.

[12]This in itself is noteworthy, as analytically tractable models, such as the Gordon (1962) growth model, typically imply constant or near-constant dividend yields and very little return volatility. In contrast, dividend yields observed in practice vary considerably over time and are strongly autocorrelated, and returns exhibit considerable volatility.

III Model Extensions, Multivariate Analysis, and Tests

The central focus in this section is on *joint* distributions of the financial statistics that emerge from our simulations: combinations of the returns, ex post equity premia, Sharpe ratios, dividend yields, *etc.*, and tests on the value of the ex ante equity premium using these joint distributions. We focus primarily on three moments of the data: the mean ex post equity premium, the excess return volatility, and the mean dividend yield. These three moments have the advantage of being the most precisely estimated and hence most informative for the value of the ex ante equity premium. Other moments that we could have considered are either largely redundant (such as the Sharpe Ratio which is a direct function of excess returns and the excess return standard deviation), or are so imprecisely estimated (for example, the ARCH(1) or AR(1) coefficients) that they would not help sharpen our estimates of the ex ante equity premium. Of course, we also do not consider the distributions of financial variables to which we calibrate our simulations (interest rates and dividend growth rates), as the simulated mean, variance, and covariance of these variables are, by construction, identical to the corresponding moments of the actual data to which we calibrate.

Our purpose in considering joint distributions is two-fold. First, multivariate tests are used to form a tight confidence bound on the true value of the ex ante equity premium. These tests strongly reject our models if the ex ante 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, this analysis leads us to reject 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 ex ante 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. The Model 1 simulation incorporates some appealing basic features, such as parameter uncertainty and calibrated time-series models for equity premia, interest rates, and dividend growth rates. It does not, however, incorporate some features of the equity premium process that have been indicated by other researchers. One omitted feature is a gradual downward trend in the equity premium, as documented in many studies, including Jagannathan, McGrattan, and Scherbina (2000), Pástor and

Stambaugh (2001), Bansal and Lundblad (2002), and Fama and French (2002). Another is a structural break in the equity premium process over the early 1990s, as shown by Pástor and Stambaugh (2001). An increase in the growth rate of cashflows (but not dividends) to investors starting in the late 1970s, as documented by Bagwell and Shoven (1989), Fama and French (2001) and others, is also a feature that Model 1 fails to incorporate. Therefore, in this section we consider models which incorporate one, two, or all three of these features, as well as different time-series models for interest rates and equity premia. 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 2 through 8 (to be fully discussed below), we present χ^2 test statistics for the null hypothesis that the US experience during 1952 through 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.[13]

A significant test statistic, in this context, suggests that the combination of financial statistics observed for the US economy is significantly unusual compared to the collection of simulated data, leading us to reject the null hypothesis that the given model and assumed ex ante equity premium value could have generated the US data of the last half century. It is possible to reject every ex ante 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 ex ante equity premium values are not rejected for a collection of models, thwarting our efforts to provide a precise estimate of the ex ante equity premium or a small range of allowable equity premium models.

As it happens, models that ignore breaks and trends in the equity premium are rejected for

[13] The 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 the excess return volatility, we consider the *cube root* of the return variance, which is approximately normally distributed (see page 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 US data (again, see Fama and French, 2001), we also adjust the US data to reflect the increase in mean payout levels. This makes for a small difference in the mean US payout ratio and no qualitative change to our results if ignored.

virtually every value of the ex ante equity premium we consider. But for a group of sophisticated models that incorporate trends and breaks in the equity premium, we cannot reject a narrow range of ex ante equity premia, roughly between 3% and 4%. We also find that models tend to be rejected if the impact on cashflows 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 cashflow trends and breaks, and 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 I. For Model 2, an 80 basis point downward trend is incorporated in the equity premium process. For Model 3, a 50 basis point 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 basis points, 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 basis points per year. These one-at-a-time feature additions help us evaluate if one or another feature documented in the literature can markedly improve model performance over the simple base model.

Panel A of Figure 2 and Panel A of Figure 3 display plots of the value of joint χ^2 tests on three moments of the data, the mean ex post equity premium, the excess return volatility, and the mean dividend yield, for Models 1 though 4, and shows how the test statistic varies as the ex ante equity premium varies from 2.25% to 8% in increments as small as an eighth of a percent toward the lower end of that range. Panels B through D of Figures 2 and 3 display the univariate Student t-test statistics for each of these three moments of the data, again showing how the test statistic varies with the assumed value of the ex ante equity premium. The values of the ex ante equity premia indicated on the horizontal axis represent the *ending* values of the ex ante equity premium in each set of simulations. For models which incorporate a downward trend or a structural break in the equity premium, the ending value of the ex ante equity premium differs from the starting value.

So, for instance, Model 2 has a starting ex ante equity premium that is 80 basis points higher than that displayed in Figure 2, as Model 2 has an 80 basis point trend downward in the ex ante equity premium. For Model 1 the value of the ex ante 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.

Figures 2 and 3 go about here.

Consider now specifically Panel A of Figures 2 and 3. (Note that we use a log scale for the vertical axis of the plots in Panel A of Figures 2 through 8 for clarity of presentation. Note as well that we postpone further discussion of Panels B through D until after we have introduced results for all the models, 1 through 12.) On the basis of Panel A of Figures 2 and 3, we see that 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 ex ante equity premium within about 25 basis points of 4%. Models 1-3, in contrast, are rejected at the 10% level for *every* ex ante equity premium value. If we allow fairly substantial departures of the S&P 500 data from the expected distribution, say 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 premia 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 ex ante equity premium.

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

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

We turn now to joint tests based on Models 5 though 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 CAPM (detailed in Appendix 1). These models also permit trends and/or breaks in the equity premium and dividend growth rate processes two or more at-a-time and incorporate alternative time-series models for the interest rate and the equity premium processes. Models 1 through 4 demonstrate that it is not sufficient to model the equity premium as an autoregressive time-varying process, and that one-at-a-time augmentation with trends or breaks in the equity premium process is also not sufficient, though the augmentations do lead to improvements over the base model in our ability to match sample moments from the US experience of the last half century. Models 5 through 10 allow us to explore questions like: do we need a conditionally time-varying equity premium model built on the Merton conditional CAPM 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 in our base model?

Model 5 is the base model, Model 1, augmented to include an 80 basis point gradual downward trend in the equity premium and a 100 basis point gradual upward trend in the dividend growth rate. Model 6 is the base model adjusted to incorporate a 30 basis point gradual downward trend in the equity premium, a 50 basis point abrupt decline in the equity premium, and a 100 basis point gradual upward trend in the dividend growth rate. Model 7 is the best model as indicated by the Bayesian Information Criterion (BIC),[14] augmenting the equity premium process with a 30 basis point gradual downward trend and a 50 basis point abrupt decline and adding a 100 basis point gradual upward trend in the dividend growth rate. Model 8 takes the second-best BIC model

[14]For Models 7 and 8 we employ the 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 EGARCH models).

and incorporates a 30 basis point gradual downward trend in the equity premium, a 50 basis point abrupt decline in the equity premium, and a 100 basis point gradual upward trend in the dividend growth rate. Model 9 is the base model adjusted to incorporate a 30 basis point gradual downward trend in the equity premium and a 50 basis point abrupt decline in the equity premium. Model 10 has the equity premium model following a deterministic downward trend with a 50 basis point structural break, interest rates following an AR(1), and dividend growth rates following 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 non-dividend cashflows to investors (such as share repurchases) starting in the late 1970s, we believe Models 6, 7, and 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 Panel A of Figures 4, 5, and 6 we present plots of the χ^2 test statistics on three moments of the data, the mean ex post equity premium, the excess return volatility, and the mean dividend yield. Again, we consider Panels B through D later. We see in Panel A of Figures 4 and 5 that for Models 5 through 8 we cannot reject a range of ex ante equity premium values at the 5% level. These models produce test statistics that drop well below even the 10% critical value (recall that Panel A's scale is logarithmic, and thus compressed). These models all embed the increased cashflow feature and either an eighty basis point downward trend in the equity premium, or both a break and a trend in the equity premium, adding to an eighty basis point decline over the last half century. The range of ex ante equity premia 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 less than 75 basis points at the 10% level. The range is slightly wider for Models 5 and 6, roughly 75 to 100 basis points. In each case, the ex ante 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 which exclude the cashflow increase, Models 9 and 10, displayed in Figure 6, we see that we can reject at the 10% level all ex ante equity premium values. Model 9 is best compared to Model 6, as it is equivalent to Model 6 with the sole difference of excluding the cashflow increase. We see from Panel A of Figures 4 and 6 that excluding the cashflow 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 6 (recall that the scale is compressed in Panel A as we use a log scale). Model 10 is identical to Model 9 apart from the sole difference that Model 10 excludes the Merton CAPM conditionally-varying equity premium process. Exclusion of this conditional time variation (modeled as a first order autoregressive process) worsens the ability of the model to match moments to the US experience at every value of the ex ante equity premium. The difference in performance leads us to reject a model excluding a conditionally-varying equity premium.

Figures 4, 5, and 6 go about here.

On the basis of our most plausible models, Models 6, 7, and 8, we can conservatively conclude that the ex ante equity premium is within 50 basis points 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 Merton (1980) CAPM 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 cashflows to investors over the last 25 years also compromises our ability to match the experience of US 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 full information maximum likelihood procedure produces parameter estimates that are jointly normally distributed, with an easily computed variance-covariance structure.

Now we consider two 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 ex ante equity premia supported by the last half century of US data. Model 11 is the base model augmented to incorporate a 30 basis point gradual downward trend in the equity premium, a 50 basis point abrupt decline in the equity premium, and a 100 basis point 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, Model 1, with no parameter uncertainty.

Figure 7 goes about here.

In Panel A of Figure 7 we present plots of the χ^2 test statistics on three moments of the data, the mean ex post equity premium, the excess return volatility, and the mean dividend yield. Again, we consider Panels B through D later. We see in Panel A that both Models 11 and 12 are rejected for all values of the ex ante 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. In each case, the ex ante equity premium that yields the minimum joint test statistic, corresponding to our estimate of π, is centered around 3%. It is apparent that parameter uncertainty is an important model feature. Ignoring parameter uncertainty leads to model rejection, even at the ex ante 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 ex ante equity premium in our joint χ^2 tests? An examination of the ranges of the ex ante equity premium consistent with the *individual* moments can shed some light on the source of the power of the joint tests. Panels B, C, and D of Figures 2 through 7 display plots of the univariate t-test statistics based on each of the variables we consider in the joint tests plotted in Panel A of these figures. Panel B of each figure plots t-test statistics on the ex post equity premium, Panel C of each figure plots t-test statistics on the excess return volatility, and Panel D of each figure plots t-test statistics on the price-dividend ratio.

Consider first Panel B of Figures 2 through 7. Virtually all of the models have a minimum t-test statistic at a point that is associated with an ex ante equity premium close to 6%.[15] Because our method involves minimizing the distance between the ex post equity premium based on the actual 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 ex ante equity premium close to 6%. The t-test on the mean ex post equity premium rises linearly as the ex ante equity premium setting departs from 6% for each model, but does not typically reject ex ante equity premium values at the 10% level until they deviate quite far from the ex ante value at which the minimum t-test is observed. For example, in Panel B of Figure 4 the ending ex ante 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 actual S&P 500 data.

The t-tests on the excess return volatility, presented in Panel C of Figures 2 through 7, indicate that lower ex ante equity premium values lead to models that are better able to match the S&P 500 experience of volatile returns.[16] Note that as the ex ante equity premium decreases, the volatility of returns *increases*, so high ex ante equity premia lead to simulated return volatilities that are much lower than the actual S&P 500 return volatility we have witnessed over the last half century. The test statistic, however, rises slowly as the ex ante equity premium grows larger, in contrast to the joint test statistics plotted in Panel A of Figures 2 through 7, in which the χ^2 test statistic

[15] Recall that the ex ante 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. Accordingly, for this model we observe (in Panel B of Figure 7) a minimum t-test at an *ending* value of the ex ante equity premium which is below the 6% *average* equity premium. The coarseness of the grid of ex ante equity premium values around 6% prevents this feature from being more obvious for some of the other models.

[16] The intuition behind this result is easiest to see by making reference to the Gordon (1962) constant dividend growth model, shown above in Equation 9. 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 prices changes, *i.e.* volatility.

rises *sharply* as the ex ante equity premium grows larger (recall that the Panel A vertical axis has a compressed log scale in Figures 2 through 7). Given these contrasting patterns, the return volatility moment is unlikely, *by itself,* to be causing the sharply rising joint test statistic.

Consider now the t-test statistics on the price-dividend ratio, plotted in Panel D of Figures 2 through 7. Notice that in all cases the t-test on the price-dividend ratio jumps up sharply as the ex ante equity premium rises above 3%. Thus the sharply increasing χ^2 statistics we saw in Panel A of the three 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 premia above 4% that the dividend yield also leads us to. In terms of the three moments we have considered in the joint χ^2 and univariate t-test statistics, it is evident that the upper range of ex ante equity premia consistent with the experience of the last half century in the US is limited by the high average S&P 500 price-dividend ratio (or equivalently, the low average S&P 500 dividend yield) 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 ex ante 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 US.

D.1 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 US, we perform two checks. First, in Models 4 through 8 we incorporate higher dividends and dividend growth rates than observed in US corporate dividends. This is to adjust for the practice, adopted widely beginning in the late 1970s, of US firms delivering cashflows to investors in ways (such as share repurchases) which are not recorded as corporate dividends. As we previously reported, Models 4 through 8 (the models that incorporate higher cashflows to investors than recorded by S&P 500 dividend payments, *i.e.*, the models that use cashflows *including* share repurchases) are best able to account for the observed US data. Reassuringly, the estimate of the equity premium emerging from Models 4 through 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 (unreported) joint tests confirm two 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 ex ante equity premium above 4%. Essentially, this indicates that using two moments of the data (excluding the price-dividend ratio) rather than all three makes it more difficult to identify the minimum test statistic value and thus more difficult to identify our estimate of the ex ante 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 Panel A of Figures 2 through 7. Second, and most importantly, the minimum test statistic is still typically achieved for models with an ex ante equity premium value between 3% and 4%. For some of the models, the minimum test statistic is 25 or 50 basis points lower than that found when basing joint tests on the full set of three moments. For a few models, the minimum test statistic is 25 or 50 basis points higher. Again Models 1 through 3 are rejected for every value of the ex ante equity premium, and again for Models 4 through 8 the range of ex ante equity premia that are not rejected is narrow.

E 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 *mis*specification) by looking at a variety of models for interest rates, dividend growth rates, and equity premium, 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 US over the last half century leads us to the conclusion that the range of true ex ante equity premia that could have generated the US experience is fairly narrow, under 100 basis points, centered roughly on 3.5%. We have not yet addressed, however, 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 ex ante equity premia have employed the observed return volatility over the last half century in the US 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 now consider the joint χ^2 test statistic based on only the mean return and the mean price-dividend ratio. Figure 8, Panel A plots the test statistics for Models 1, 2, and 3, Panel B plots the test statistics for Models 4, 5, and 6, Panel C plots the test statistics for Models 7, 8, and 9, and Panel D plots the test statistics for Models 10, 11, and 12, with a log scale for the vertical axis in all cases.

Figure 8 goes about here.

First consider results for Models 1 through 4, shown in Panels A and B of Figure 8. These are the base model with no trends or breaks, and models which incorporate only one feature (trend or break in the equity premium or dividend growth rate) at a time. We see again that Model 1 is rejected outright for every value of the ex ante equity premium, at the 10% level of significance, and we see again that adding trends or breaks, even one-at-a-time, improves performance. Now Model 2 (incorporating an 80 basis point downward trend in the equity premium) and Model 4 (incorporating the increased cashflow growth rate) are not rejected over narrow ranges at the 10% significance level. We find that Models 5, 6, 7, and 8, all incorporating trends and breaks in the equity premium and dividend growth rate processes and shown in Panels B and C of Figure 8, deliver a wide range of ex ante equity premia which cannot be rejected at any conventional level of statistical significance. We also see that Model 9 in Panel C, incorporating a trend (of 30 basis points) and a break (of 50 basis points) in the equity premium, performs similarly to Model 2, which has only a trend of 80 basis points (neither model incorporates a cashflow change). In Panel D we see Model 10 which has a deterministic equity premium with trends and breaks. This model's performance is also similar to Model 2, but slightly worse, rejected at the 10% level at every ex ante equity premium. Also in Panel D we see that Models 11 and 12, which do not incorporate parameter estimation uncertainty, are almost everywhere rejected. (In contrast to the joint test shown in Panel A of Figure 7, based on all three moments, we find that Model 11 is not rejected only for the 3% value of the ex ante equity premium.)

Overall, the value of the ex ante equity premium at which the joint test statistic is minimized (*i.e.*, our estimate of the ex ante equity premium) is not particularly affected by our having based the joint tests on two moments of the data rather than the original three, nor is our selection of plausible models for the equity premium process. Across the models, the highest estimate of the ex ante equity premium is roughly 4% (for Model 4) and the lowest is 3% (for Models 11 and 12). With the joint tests based on two moments, all models support (*i.e.*, do not reject) broader ranges of the ex ante equity premium, with the range widest for Models 4 through 8 (now spanning roughly 200 basis points for any given model, from ex ante 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 ex ante equity premia is consistent with a decline in the power of our joint test, presumably from omitting an important moment of the data, the return volatility. The widening of the range of plausible ex ante equity premia 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 US 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 ex ante equity premium.

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

F Bootstrapped Test Statistics

Up to this point, all of our test statistics have relied on asymptotic distribution theory for critical values. The asymptotic distributions should be reliable both because we are looking at averages over independent events (our simulations are by construction independent) and because we have many simulations over which to average (2,000). Nonetheless, it is straightforward to use our simulated test statistics to bootstrap the distribution of the test statistics, thus we do so. While use of the bootstrap produces small quantitative changes to our results, our main findings remain unchanged. The best estimate of the mean ex ante equity premium and the range of plausible ex ante equity premia and equity premium models do not budge.

IV 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 actually receive ex post, after they purchase the stock and hold it over some period of time during which random economic shocks impact prices. US stocks have historically returned roughly 6% more than risk-free debt. Ex post estimates provided by recent papers suggest the US equity premium may be falling in recent years. However, all of these estimates are imprecise, and there is little consensus emerging about the true value of the ex ante equity premium. The imprecision and lack of consensus both 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 resolution of the equity premium puzzle and makes it difficult to determine if the equity premium changes over time.

In order to determine the most plausible value of the ex ante 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 in order to improve the estimate of the mean ex ante equity premium follows recent work on the equity premium which has also sought improvements through the use fundamental information like the dividend and earnings yields (Fama and French, 2002, and Jagannathan, McGrattan, and Scherbina, 2000), higher-order moments of the excess return distribution (Maheu and McCurdy, 2007) and return volatility and price movement directions (Pástor and Stambaugh, 2001).

Our central insight is that the knowledge that a low dividend yield, high ex post equity premium, high return volatility, and high Sharpe ratio all occurred together over the last five decades tells us something about the mean ex ante equity premium and the likelihood that the equity premium is time-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 ex post equity premium. 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 employ the simulated method of moments technique and build on the dividend discounting method of fundamental valuation of Donaldson and Kamstra (1996) to estimate the ex ante equity premium. We reject as inconsistent with the US experience all but a narrow range of values of the mean ex ante equity premium and all but a small number equity premium time-series models. We do so while incorporating model estimation uncertainty and allowing for investor uncertainty about the true state of the world. The range of ex ante equity premia 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 – that is, consistent with the experience of the US over the last half century – incorporate substantial autocorrelation, a structural break, and/or a gradual downward trend in the equity premium process. For these models, the range of ex ante equity premia supported by our tests is very narrow, plus or minus 50 basis points around 3.5%. All together, our tests strongly support the notion that the equity premium process over the last half century in the US was very unlikely to have been constant, was likely to have demonstrated at least one sharp downward break, and was likely to have demonstrated a gradual downward trend.

References

Bagwell, L., Shoven, J., 1989, Cash distributions to shareholders, *Journal of Economic Perspectives* 3, 129-149.

Bansal, Ravi and Amir Yaron, 2004, Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles, *Journal of Finance* 59(4), 1481-1509.

Bansal, Ravi and Christian Lundblad, 2002, Market efficiency, asset returns, and the size of the risk premium in global equity markets, *Journal of Econometrics* 109, 195-237.

Bekaert, Geert and Campbell R. Harvey, 1995, Time-varying world market integration, *Journal of Finance* 50, 403-444.

Campbell, John Y., 1991, A Variance Decomposition for Stock Returns, *The Economic Journal* 101, 157-179.

Campbell, John Y. and Robert J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195-228.

Campbell, John. Y., and Robert. J. Shiller, 1998, Valuation Ratios and the Long-Run Stock Market Outlook, *Journal of Portfolio Management* 24(6), 11-26.

Carrasco, Marine and Jean-Pierre Florens, 2002, Simulation-based method of moments and efficiency, *Journal of Business and Economic Statistics* 20(4), 482-492.

Claus, James and Jacob Thomas, 2001, Equity Premia as Low as Three Percent? Evidence from Analysts' Earnings Forecasts for Domestic and International Stock Markets, *Journal of Finance* 56(5), 1629-1666.

Cochrane, John H., 2001 *Asset Pricing,* Princeton NJ, Princeton University Press.

Corradi, Valentina and Norman R. Swanson, 2005, Bootstrap specification tests for diffusion processes, *Journal of Econometrics* 124(1), 117-148.

Diamond, Peter A., 2000, What stock market returns to expect for the future? *Social Security Bulletin* 63(2), 38-52.

Dimson, Elroy, Paul Marsh, and Mike Staunton, 2007, The Worldwide Equity Premium: A Smaller Puzzle, in Rajnish Mehra, ed.: *Handbook of Investments: Equity Risk Premium*, Amsterdam: Elsevier.

Donaldson, R. Glen and Mark J. Kamstra, 1996, A New Dividend Forecasting Procedure That Rejects Bubbles in Asset Prices, *Review of Financial Studies* 9, 333-383.

Duffie, Darrell and Kenneth J. Singleton, 1993, Simulated Moments Estimation of Markov Models of Asset Prices, *Econometrica* 61(4), 929-952.

Fama, Eugene F. and Kenneth R. French, 1997, Industry Costs of Equity, *Journal of Financial Economics* 43, 153-193.

Fama, Eugene F. and Kenneth R. French, 2001, Disappearing dividends: Changing firm characteristics or lower propensity to pay, *Journal of Financial Economics* 60, 3-43.

Fama, Eugene F. and Kenneth R. French, 2002, The Equity Premium, *Journal of Finance* 57 (2), 637-659.

Graham, John, and Campbell Harvey, 2005, The Long-Run Equity Risk Premium, *Finance Research Letters* 2, 185-194.

Gregory, Allan W. and Gregor W. Smith, 1991, Calibration as Testing: Inference in Simulated Macroeconomic Models. *Journal of Business and Economic Statistics* 9(3), 297-303.

Gordon, Myron, 1962, *The Investment, Financing and Valuation of the Corporation*, Homewood, IL: Irwin.

Grullon, Gustavo and Roni Michaely, 2002, Dividends, Share Repurchases, and the Substitution Hypothesis, *Journal of Finance* 57(4), 1649-1684.

Hannan, E.J., 1980, The Estimation of the Order of an ARMA Process, *The Annals of Statistics* 8, 1071-1081.

Hodrick, Robert J., 1992, Dividend Yields and Expected Stock Returns: Alternative Procedures for inference and Measurement, *Review of Financial Studies* 5(3), 357-386.

Hull, John, 1993, *Options, Futures and Other Derivative Securities,* Englewood Cliffs, NJ: Prentice-Hall.

Jagannathan, Ravi, Ellen R. McGrattan, and Anna Scherbina, 2000, The Declining US Equity Premium, NBER Working Paper W8172, *Federal Reserve Bank of Minneapolis Quarterly Review* 24, 3-19.

Jain, Pankaj K. 2005, Financial Market Design and the Equity Premium: Electronic versus Floor Trading, *Journal of Finance* 60(6), p2955-2985

Jorion, Philippe and William N. Goetzmann, 1999, Global Stock Markets in the Twentieth Century, *Journal of Finance* 54(3), 953-980.

Kendall, Maurice and Alan Stuart, 1977, *The Advanced Theory of Statistics,* Volume 1, 4th edition, New York: MacMillan Publishing.

Kleidon, Allan W., 1986, Variance Bound Tests and Stock Price Valuation Models, *Journal of Political Economy* 94, 953-1001.

Kocherlakota, Narayana R., 1996, The Equity Premium: It's Still a Puzzle, *Journal of Economic Literature* 34(1), 42-71.

Maheu, John M. and Thomas H. McCurdy, 2007, How useful are historical data for forecasting the long-run equity return distribution? forthcoming, *Journal of Business and Economic Statistics*.

Mankiw, N. Gregory, David H. Romer, and Matthew D. Shapiro, 1991, Stock Market Forecastability and Volatility: A Statistical Appraisal, *Review of Economic Studies* 58(3), 455-477.

McFadden, Daniel, 1989, A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration, *Econometrica* 47, 995-1026.

Mehra, Rajnish, 2003, The Equity Premium: Why Is It A Puzzle? *Financial Analysts Journal* 59(1) 54-69.

Mehra, Rajnish and Edward C. Prescott, 1985, The Equity Premium: A Puzzle, *Journal of Monetary Economics* 15, 145-161.

Mehra, Rajnish and Edward C. Prescott, 2003, The Equity Premium in Retrospect, in Constantinides, G.M., M. Harris, and R. Stulz, eds.: *Handbook of the Economics of Finance* (Elsevier).

Merton, Robert C., 1980, On Estimating the Expected Return on the Market: An Exploratory Investigation, *Journal of Financial Economics* 8, 323-361.

Nelson, Daniel B., 1991, Conditional Heteroskedasticity in Asset Returns: A New Approach, *Econometrica* 59, 347-370.

Newey, Whitney K. and Kenneth D. West, 1987, A Simple, Positive, Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* 55, 703-708.

Pakes, Ariel and David Pollard, 1989, Simulation and the Asymptotics of Optimization Estimators, *Econometrica* 57, 1027-1057.

Pástor, Ľuboš and Robert Stambaugh, 2001, The Equity Premium and Structural Breaks, *Journal of Finance* 56(1), 1207-1239.

Rozeff, Michael S., 1984, Dividend yields are equity risk premiums, *Journal of Portfolio Management* 11(1), 68-75.

Scott, Louis O., 1985, The Present Value Model of Stock Prices: Regression Tests and Monte Carlo Results, *The Review of Economics and Statistics* 67(4), 599-605.

Siegel, Jeremy J., 1999, The shrinking equity premium, *Journal of Portfolio Management*, 26(1), 10-17.

Shiller, Robert J., 1984, Stock Prices and Social Dynamics, *Brookings Papers on Economic Activity*, 457-498.

Siegel, Jeremy J. and Richard H. Thaler, 1997, Anomalies: The Equity Premium Puzzle, *Journal of Economic Perspectives* 11(1), 191-200.

Timmermann, Allen G., 1993, How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices, *The Quarterly Journal of Economics* 108(4), 1135-1145.

Timmermann, Allen G., 1995, Tests of Present Value Models with a Time-Varying Discount Factor, *Journal of Applied Econometrics* 10(1), 17-31.

Welch, I., 2000, Views of Financial Economists on the Equity Premium and on Professional Controversies, *Journal of Business* 73(4), 501-537.

West, Kenneth D., 1988a, Bubbles, Fads and Stock Price Volatility Tests: A Partial Evaluation, *Journal of Finance* 43, 639-656.

West, Kenneth D., 1988b, Dividend Innovations and Stock Price Volatility, *Econometrica* 56, 37-61.

White, Halbert, 1984, *Asymptotic Theory for Econometricians,* Orlando: Academic Press.

Appendices

Appendix 1: Models for Generating Data

In creating distributions of financial variables modeled on the US economy, we must generate the fundamental factors that drive asset prices: dividends and discount rates (where the discount rate is defined as the risk-free rate plus a possibly time-varying equity premium). Thus we must specify time-series models for dividend growth, interest rates, and ex ante equity premia so that our Monte Carlo simulations will generate dividends and discount rates that share key features with observed S&P 500 dividends and US discount rates. We consider a range of models to generate data in our simulations, as outlined in Table I. Each model incorporates specific characteristics that define the way we generate interest rates and dividend growth rates, and each model makes specific assumptions about the way the ex ante equity premium evolves over time, if indeed it does evolve over time. In providing further information about these defining aspects of our models, we consider each model feature from Table I in turn, starting with the time-series processes for interest rates, dividend growth rates, and the ex ante equity premium.

A1.1 Processes for the Interest Rate, Dividend Growth Rate and the Ex Ante Equity Premium

The interest rate and dividend growth rate series we generate are calibrated to the time-series properties of data observed in the US over the period 1952 to 2004. We considered the ability of various time-series models to eliminate residual autocorrelation and ARCH (evaluated with 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 one 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 non-linear 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, page 408), perform well in capturing the time-series properties of observed interest rates. We also find the AR(1) and ARMA(1,1) specifications perform comparably to one another, 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 log-normality 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 the log-normal. Thus we model the log of 1 plus the dividend growth rate, and we find that both a MA(1) and an AR(1) specification fit the data well, removing evidence of residual autocorrelation and ARCH at five 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 log-normality of dividend growth rates) do not reject the null of normality.

Most of our models incorporate an ex ante equity premium that follows an ARMA process emerging from Merton's (1980) conditional CAPM. Merton's 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 i^{th} asset,

$$E_t(r_{i,t}) = \lambda \; cov_{t-1}(r_{i,t} r_{m,t}), \tag{10}$$

where $r_{i,t}$ are excess returns on the asset, $r_{m,t}$ are excess returns on the market portfolio, cov_{t-1} is the time-varying conditional covariance between excess returns on the asset and on the market portfolio, and E_t is the conditional-expectations operator incorporating information available to the market up to but not including the beginning of period t. λ is a parameter of the model, described below.

For the expected excess market return, (10) becomes

$$E_t(r_{m,t}) = \lambda \; var_{t-1}(r_{m,t}) \tag{11}$$

where var_{t-1} is the market time-varying conditional variance. Merton (1980) argues that λ in (11) 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 (11) defines a time-varying equity premium but has the equity premium varying only as a function of time-varying conditional variance. Following Bekaert and Harvey (1995), it is possible to allow λ in Equation (11) to vary over time by making it a parametric function of conditioning variables (indicated below as Z_{t-1}). The functional form Bekaert and Harvey employ (in Equation (12) of their paper) is exponential, restricting the price of risk to be positive:

$$\lambda_{t-1} = exp\left(\delta' Z_{t-1}\right). \tag{12}$$

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 $var_{t-1}(r_{m,t})$. Once again we calibrate to the S&P 500 over 1952 to 2004, estimating the following model:

$$r_{m,t} = \lambda_{t-1}\ var_{t-1}(r_{m,t}) + e_{m,t} \tag{13}$$

$$var_{t-1}(r_{m,t}) = \omega + \alpha e^2_{m,t-1} \tag{14}$$

$$\lambda_{t-1} = exp\left(\delta_0 + \delta_1 \frac{D_{t-1}}{P_{t-1}}\right). \tag{15}$$

The values of estimated parameters are $\delta_0 = -3.93$, $\delta_1 = 0.277$, $\omega = 0.0194$, and $\alpha = 0.542$. The R^2 of this model is 2.8%.

For our simulations, we model the time-series process of the ex ante time-varying equity premium (denoted π_t) by using the excess return as a proxy for the equity premium:

$$\hat{\pi}_t = \hat{\lambda}_{t-1} v\hat{a}r_{t-1}(r_{m,t}), \tag{16}$$

where $\hat{\lambda}_{t-1} = exp\left(-3.93 + 0.277\frac{D_{t-1}}{P_{t-1}}\right)$, $v\hat{a}r_{t-1}(r_{m,t}) = 0.0194 + 0.542\hat{e}^2_{m,t-1}$, and $\hat{e}_{m,t-1} = r_{m,t-1} - \hat{\pi}_{t-1}$. The time-varying 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,[17] so that when the equity premium is perturbated it reverts to its mean slowly. This permits slightly more volatile returns in our simulations than would otherwise be the case. The best way to see the impact of this slow mean reversion of the equity premium on our simulations is to compare Models 9 and 10. Model 9 has a conditionally time-varying equity premium (together with a trend and break in the premium) while Model 10 is identical except the equity premium does not conditionally vary. We find standard tests for normality of the error term (and hence conditional log-normality of the equity premium) show some evidence of non-normality when estimated as a single equation, but less or no evidence if estimated in a system of equations with the interest rate and dividend growth rate equations.

Hence we generate the ex ante equity premia, interest rate, and dividend growth rate series as autocorrelated series with jointly normal error terms, calibrated to the degree of autocorrelation observed in the US data. The processes we simulate also mimic the covariance structure between the residuals from the time-series models of equity premia, interest rates, and dividend growth rates as estimated using US data. We adjust the mean and the standard deviation of these log-normal 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 A.I the estimated parameters of Model 1, which incorporates an AR(1) model for interest rates (r), a MA(1) model for dividend growth rates (g), and an AR(1) model for the ex ante equity premium (π).

[17] The 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 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.

Table A.I
Estimated Parameters of Model 1

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

In Table A.I, standard errors of the estimated coefficients are shown 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%.

A1.2 Allowing a Downward Trend in the Ex Ante 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 0.80% 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, we investigate a 0.80% trend in the equity premium, and when modeling a trend with a break we limit ourselves to a 0.30% trend with an additional 50 basis point break, as discussed below. This is accomplished in conjunction with setting the ex ante equity premium to follow an AR(1) process.

A1.3 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 two centuries. They find fairly strong support for there having been a structural break over the 1990s which 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 ex ante equity premium). Spreading the drop in the premium across several years serves only to narrow the range of ex ante equity premium consistent with the US 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 ex ante equity premium. Thus we incorporate an abrupt 50 basis point drop in the equity premium in some of the models we consider. We time the drop to coincide with 1990, 39 years into our simulation period. This feature of the equity premium process can be accomplished with or without incorporating other features discussed above.

A1.4 Allowing for Sampling Variability in Generating Parameters

Our experiments are motivated by the large sampling variability of the ex post equity premium, but when we produce our simulations we have to first estimate the parameter values for the time-series models of dividend growth rates, interest rates, and ex ante equity premia. 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 the ex ante equity premium equation) 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 considerations[18] and data consistency checks.[19] This process accounts for possible variability in the true state of the world that generates dividends, interest rates, and ex ante equity premia.

To illustrate, for Model 1 reported in Table A.I,

[18] The 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 zero.

[19] The parameters must generate mean interest rates, dividend growth rates, and ex post equity premia that lie within three standard deviations of the US data sample mean. Also, the limiting price-dividend ratio must be within 50 standard deviations of the mean US 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 US data.

$$
\begin{array}{lcllc}
log(r_t) & = & \alpha_r & +\rho_r \, log(r_{t-1}) & +\epsilon_{r,t} \\
log(1+g_t) & = & \alpha_g & +\theta_g \, \epsilon_{g,t-1} & +\epsilon_{g,t} \\
log(\hat{\pi}_t) & = & \alpha_\pi & +\rho_\pi \, log(\hat{\pi}_{t-1}) & +\epsilon_{\pi,t},
\end{array}
$$

the estimated covariance matrix of the parameter estimates is shown in Table A.II.

Table A.II
Estimated Covariance Matrix for Model 1 Parameters

	α_r	ρ_r	α_g	θ_g	α_π	ρ_π
α_r	0.068705	0.022307	-.000051933	.000226443	-0.012165	-0.003511
ρ_r	0.022307	0.007436	-.000040346	.000114831	-0.004730	-0.001401
α_g	-0.000052	-0.000040	0.000039674	.000025651	0.000153	0.000031
θ_g	0.000226	0.000115	0.000025651	.007086714	0.001699	0.000454
α_π	-0.012165	-0.004730	0.000153376	.001699151	0.052664	0.015791
ρ_π	-0.003511	-0.001401	0.000031495	.000453874	0.015791	0.004844

The top-left element of Table A.II, 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. The *estimated covariance matrix* of the equation residual variances is shown in Table A.III. (The variances themselves are reported in Section A1.1, as are the parameter estimates of the mean.)

Table A.III
Estimated Covariance Matrix of Model 1 Residual Variances

	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	$1.9729\cdot10^{-6}$	$-8.351\cdot10^{-7}$	$-1.902\cdot10^{-7}$	$-1.564\cdot10^{-6}$	$-1.69\cdot10^{6}$
$\epsilon_r\epsilon_g$	$1.9729\cdot10^{-6}$	$8.5163\cdot10^{-7}$	$1.0437\cdot10^{-6}$	$4.3066\cdot10^{-8}$	$-1.602\cdot10^{-7}$	$9.1448\cdot10^{-7}$
$\epsilon_r\epsilon_\pi$	$-8.351\cdot10^{-7}$	$1.0437\cdot10^{-6}$	0.0000797	$1.8827\cdot10^{-7}$	$5.001\cdot10^{-6}$	-0.000044
ϵ_g^2	$-1.902\cdot10^{-7}$	$4.3066\cdot10^{-8}$	$1.8827\cdot10^{-7}$	$4.8337\cdot10^{-8}$	$9.6885\cdot10^{-8}$	$1.3458\cdot10^{-6}$
$\epsilon_g\epsilon_\pi$	$-1.564\cdot10^{-6}$	$-1.602\cdot10^{-7}$	$5.001\cdot10^{-6}$	$9.6885\cdot10^{-8}$	$3.5567\cdot10^{-6}$	0.0000203
ϵ_π^2	$-1.69\cdot10^{-6}$	$9.1448\cdot10^{-7}$	-0.000044	$1.3458\cdot10^{-6}$	0.0000203	0.0005009

The top-left element, equal to 0.0000944, is the variance of ϵ_r^2. The entry below the top-left element, equal to $-1.9729\cdot10^{-6}$, is the covariance between the estimate of ϵ_r^2 and the product of ϵ_r and ϵ_g, and so on.

Exploiting block diagonality of the parameters of the mean and variance, and asymptotic normality of all the estimated parameters, we generate two sets of normally distributed random variables.

Each set is independent of the other, the first set of six having the covariance matrix from Table A.II with means equal to the parameter estimates listed in Table A.I, and the second set of six having the covariance matrix from Table A.III, with means equal to the equation residual covariances listed in Section A1.1. This set of 12 random variables is then used to simulate interest rates, dividend growth rates, and equity premia, subject to the consistency checks footnoted earlier.

A1.5 Allowing for Disappearing Dividends

An issue with our calibration to dividends is the impact of declining dividend payments in the US. This phenomenon is a result of a practice adopted widely beginning in the late 1970s, whereby US firms have been increasingly delivering cashflows to investors in ways not recorded as corporate dividends, such as share repurchases. Fama and French (2001) document the widespread decline of regular dividend payments starting in 1978, consistent with evidence provided by Bagwell and Shoven (1989) and others. Fama and French find evidence that the disappearance of dividends is in part due to an increase in the inflow of new listing to US stock exchanges, representing mostly young companies with the characteristics of firms that would not be expected to pay dividends, and in part due to a decline in the propensity of firms to pay dividends. Fama and French find only a small decline in the probability to pay dividends among the firms that we calibrate to, those in the S&P 500 index.

Consistent with Fama and French, we find no evidence of a break in our data on dividend growth rates. Though 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 dividends declining in any absolute sense. 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 have averaged 4.9% with an annual standard deviation of 3.9%, and from 1979 to 2000 the dividend growth rates have 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 two periods. Consistent with this stability of dividend growth pre- and post-1978 and Bagwell and Shoven's 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 cashflows 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 basis points, to proxy for the increase in total cashflows to investors. That is, we increase the growth rate of dividends by 5 basis points 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 basis points, mimicking the proportional increase in earnings growth rates.

Appendix 2: Further Details on the Simulations

A2.1 Fundamentals

We define P_t as a stock's beginning-of-period-t price and E_t as the expectations operator conditional on information available up to but not including the beginning of period t. The discount rate (r_t, which equals the risk-free rate plus the equity premium) is the rate investors use to discount payments received during period t (*i.e.*, from the beginning of period t to the beginning of period $t+1$). Recall that investor rationality requires that the time t market price of a stock, which will pay a dividend D_{t+1} one period later and then sell for P_{t+1}, satisfy Equation (3):

$$P_t = E_t\left\{\frac{P_{t+1} + D_{t+1}}{1 + r_t}\right\}. \tag{3}$$

Invoking the standard transversality condition that the expected present value of the stock price P_{t+i} falls to zero as i goes to infinity, and defining the growth rate of dividends during period t as $g_t \equiv (D_{t+1} - D_t)/D_t$, allows us rewrite Equation (3) as:

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

One 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 actual 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, $\frac{1+([g_t-I_t]/[1+I_t])}{1+([r_t-I_t]/[1+I_t])} = \frac{(1+g_t)}{(1+r_t)}$, 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.

Properties of prices and returns produced by Equation (5) depend in important ways on the modeling of the dynamics of the dividend growth, interest rate, and equity premium processes. For instance, the stock price would equal a constant multiple of the dividend level and returns would be very smooth over time if dividend growth and interest rates were set equal to constants plus independent innovations. However, using models that capture the serial dependence of dividend growth rates, interest rates, and equity premia observed in the data, as we do, would typically lead to time-varying price-dividend ratios and variable returns of the sort we observe in observed stock market data.

A2.2 Numerical Simulation

We now provide details on the numerical simulation which comprises Step 4 of the 5-step procedure outlined in Section I above. That is, we detail for the n^{th} economy the formation of the prices (P_t^n), returns (R_t^n), ex post equity premia ($\hat{\pi}^n$), *etc.* (where $n = 1, \cdots, N$ and $t = 1, \cdots, T$), given dividends, dividend growth rates, risk-free interest rates, and the equity premium of the n^{th} economy: D_t^n, g_{t-1}^n, and $r_{t-1}^n = r_{f,t-1}^n + \pi$.[20] For simplicity, we illustrate our methodology by assuming fixed parameters (no parameter uncertainty), a constant ex ante equity premium, and an AR(1) model for interest rates. Further, to illustrate the procedure required for a moving average error model, we assume a MA(1) process for dividend growth rates. Relaxing these assumptions (the assumptions to incorporate parameter uncertainty, ARMA(1,1) processes for interest rates and dividend growth rates, and a time-varying equity premium) complicates the procedure outlined below only slightly. Note that in our actual simulations we set the initial dividend growth rate and

[20]We 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.

interest rate to their unconditional means, innovations to zero, and dividends to \$1, then simulate the economies out for 50 periods. At period 51 we start our calculation of market prices, returns, *etc.* (to avoid contaminating the simulations with the initial conditions). For simplicity, we do not include this detail in the description below but for concreteness we describe a similar prototypical simulation.

In terms of timing and information, recall that P_t^n is the stock's beginning-of-period-t price, r_t^n is the rate used to discount payments received during period t and is known at the beginning of period t, D_t^n is paid at the beginning of period t, g_t^n is defined as $(D_{t+1}^n - D_t^n)/D_t^n$ and is not known at the beginning of period t since it depends on D_{t+1}^n, and $E_t\{\cdot\}$ is the conditional expectation operator, with the conditioning information being the set of information available to investors up to but not including the beginning of period t. Finally, recall Equation (5), rewritten to correspond to the n^{th} economy:

$$P_t^n = D_t^n E_t \left\{ \sum_{i=0}^{\infty} \left(\Pi_{k=0}^{i} \left[\frac{1+g_{t+k}^n}{1+r_{t+k}^n} \right] \right) \right\}. \tag{17}$$

Returns are constructed as $R_t^n = (P_{t+1}^n + D_{t+1}^n - P_t^n)/P_t^n$, and $\hat{\pi}^n = \overline{R}^n - \overline{r}_f^n$ where $\overline{R}^n = \frac{1}{T}\sum_{t=1}^{T} R_t^n$ and $\overline{r}_f^n = \frac{1}{T}\sum_{t=1}^{T} r_{f,t}^n$.

Based on Equation (17), we generate prices by generating a multitude of possible streams of dividends and discount rates, present-value discounting the dividends with the discount rates, and averaging the results, *i.e.*, by conducting a Monte Carlo integration.[21] Hence we produce prices (P_t^n), returns (R_t^n), ex post equity premia ($\hat{\pi}^n$), and a myriad of other financial quantities, utilizing only dividend growth rates and discount rates. The *exact* procedure by which we conduct this numerical simulation is described below and summarized in Figure A.1. (These steps, labeled Steps 4A through 4C, collectively constitute Step 4 of the 5-step procedure outlined in Section I above.)

[21] According to Equation (17), the stream of dividends and discount rates should be infinitely long, however truncating the stream at a sufficiently distant point in time denoted I leads to a very small approximation error. We discuss this point more fully below.

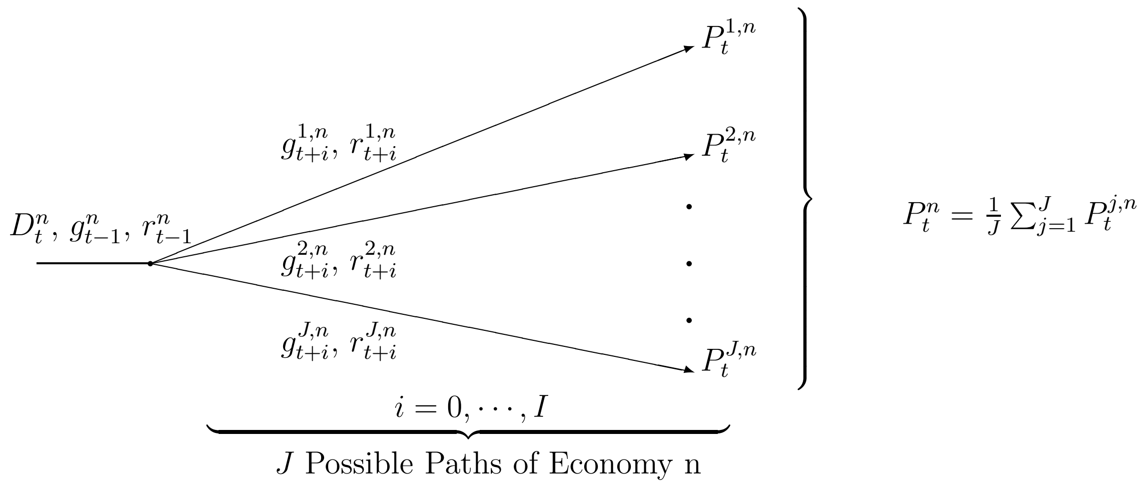

Figure A.1 *Diagram of a Simple Market Price Calculation for the t^{th} Observation of the n^{th} Economy (Steps 4A and 4B)*

Step 4A: In forming P^n_t, the most recent fundamental information available to an investor would be g^n_{t-1}, D^n_t, and r^n_{t-1}. Thus g^n_{t-1}, D^n_t, and r^n_{t-1} must be generated directly in our simulations, whereas P^n_t is calculated based on these g, D, and r. The objective of Steps 4A(i)-(iii) outlined below 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 US interest rates. In terms of Figure A.1, Step 4A forms g^n_{t-1}, D^n_t, and r^n_{t-1} only.

Step 4A(i): Note that since, as described above, the logarithm of one plus the dividend growth rate is modeled as a MA(1) process, $log(1 + g^n_t)$ is a function of only innovations, labeled ϵ^n_g. Note also that since the logarithm of the interest rate is modeled as an AR(1) process, $log(r^n_{f,t})$ is a function of $log(r^n_{f,t-1})$ and an innovation labeled ϵ^n_r. Set the initial dividend, D^n_1, equal to the total S&P 500 dividend value for 1951 (observed at the end of 1951), and the lagged innovation of the logarithm of the dividend growth rates $\epsilon^n_{g,0}$ to 0. To match the real-world interest rate data, set $log(r^n_{f,0}) = -2.90$ (the mean value of log interest rates required to produce interest rates matching the mean of observed T-bill rates). Then generate two independent standard normal random numbers, η^n_1 and ν^n_1 (note that the subscript on these random numbers indicates time, t), and form two correlated random variables, $\epsilon^n_{r,1} = 0.319(0.25\eta^n_1 + (1 - .25^2)^{.5}\nu^n_1)$ and $\epsilon^n_{g,1} = 0.0311\eta^n_1$. 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 $log(r_{f,1}^n) = -0.35+0.88log(r_{f,0}^n)+\epsilon_{r,1}^n$ to match the parameters estimated on the S&P 500 index data 1952-2004 of these models (using Full Information Maximum Likelihood).[22] Also form $D_2^n = D_1^n(1+g_1^n)$.

Step 4A(ii): Produce two correlated normal random variables, $\epsilon_{r,2}^n$ and $\epsilon_{g,2}^n$ as in Step 4A(i) above, and conditioning on $\epsilon_{g,1}^n$ and $log(r_{f,1}^n)$ from Step 4A(i) produce $log(1+g_2^n) = 0.049+0.64\epsilon_{g,1}^n+\epsilon_{g,2}^n$, $log(r_{f,2}^n) = -0.35+0.88log(r_{f,1}^n)+\epsilon_{r,2}^n$, and $D_3^n = D_2^n(1+g_2^n)$.

Step 4A(iii): Repeat Step 4A(ii) to form $log(1+g_t^n)$, $log(r_{f,t}^n)$, and D_t^n for $t = 3,4,5,\cdots,T$ and for each economy $n = 1,2,3,\cdots,N$. Then calculate the dividend growth rate g_t^n and the discount rate r_t^n (which equals $r_{f,t}^n$ plus the ex ante equity premium).

Step 4B: For each time period $t = 1,2,3,\cdots,T$ and economy $n = 1,2,3,\cdots,N$ 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 $t-1$ 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 $t, t+1, t+2, \cdots$ given what is known at period $t-1$ and use these to solve the expectation of Equation (17). We use the superscript j to index the possible paths of future economies that could possibly evolve from the current state of the economy. In Step 4B(iv) below, we describe how we are able to solve for the expectation of an infinite sum using a *finite* stream of future dividends.

Step 4B(i): Set $\epsilon_{g,t-1}^{j,n} = \epsilon_{g,t-1}^n$ and $log(r_{f,t-1}^{j,n}) = log(r_{f,t-1}^n)$ for $j = 1,2,3,\cdots,J$.[23] Generate two independent standard normal random numbers, $\eta_t^{j,n}$ and $\nu_t^{j,n}$, and form two correlated random variables $\epsilon_{r,t}^{j,n} = 0.319(0.25\eta_t^{j,n}+(1-.25^2)^{.5}\nu_t^{j,n})$ and $\epsilon_{g,t}^{j,n} = 0.0311\eta_t^{j,n}$ for $j = 1,2,3,\cdots,J$.[24] These

[22]Note that by construction these parameters do not match those reported for the system reported in Appendix 1 as this system does not incorporate a time-varying equity premium.

[23]We choose J to lie between 1,000 and 100,000, as needed 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 computed the mean and variance of the price estimate. If the experiment had no simulation error, each of the price estimates would be identical. 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 and the ex ante equity premium.

[24]For our random number generation we made 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.

are the simulated innovations to the interest rate and dividend growth rate processes, respectively. Form $log(1+g_t^{j,n}) = 0.049 + 0.64\epsilon_{g,t-1}^{j,n} + \epsilon_{g,t}^{j,n}$ and $log(r_{f,t}^{j,n}) = -0.35 + 0.88log(r_{f,t-1}^{j,n}) + \epsilon_{r,t}^{j,n}$.

Step 4B(ii): Produce two correlated normal random variables $\epsilon_{r,t+1}^{j,n}$ and $\epsilon_{g,t+1}^{j,n}$ as in Step 4B(i) above, and conditioning on $\epsilon_{g,t}^{j,n}$ and $log(r_{f,t}^{j,n})$ from Step 4B(i) produce $log(1+g_{t+1}^{j,n}) = 0.049 + 0.64\epsilon_{g,t}^{j,n} + \epsilon_{g,t+1}^{j,n}$ and $log(r_{f,t+1}^{j,n}) = -0.35 + 0.88log(r_{f,t}^{j,n}) + \epsilon_{r,t+1}^{j,n}$ for $j = 1, 2, 3, \cdots, J$.

Step 4B(iii): Repeat Step 4B(ii) to form $log(1+g_{t+i}^{j,n})$ and $log(r_{t+i}^{j,n})$ for $i = 2, 3, 4, \cdots, I$, $j = 1, 2, 3, \cdots, J$, and economies $n = 1, 2, 3, \cdots, N$.

Step 4B(iv): The discounted present value of each of the individual J streams of dividends is now taken in accordance with Equation (17), 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 (17) the stream of discount rates and dividend growth rates should be infinitely long, while in our simulations we extend the stream for only a finite number of periods, I. Since the ratio of gross dividend growth rates to gross discount rates are less than unity in steady state, the individual product elements in the infinite sum in Equation (17) eventually converge to zero as I increases. (Indeed, this convergence to zero is exactly what is required for the standard transversality condition that the expected present value of the stock price P_{t+i} falls to zero as i goes to infinity.) We therefore set I large enough in our simulations so that the truncation does not materially effect our results. We find that setting $I = 1,000$ years is sufficient in all cases we studied. That is, the discounted present value of a dividend payment received 1,000 years in the future is essentially zero. Also note that the steps above are required to produce P_t^n, D_t^n, g_t^n, and r_t^n for $n = 1, \cdots, N$ and $t = 1, \cdots, T$; the intermediate terms superscripted with a j are required only to perform the numerical integration that yields P_t^N. Note that 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.

Step 4C: After performing Steps 4A(i)-(iii) and 4B(i)-(iv) for $t = 1, \cdots, T$, rolling out N independent economies for T periods, we construct the market returns for each economy, $R_t^n = (P_{t+1}^n + D_{t+1}^n - P_t^n)/P_t^n$, 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.

Table I

Characteristics of Simulated Models

Here we present the 12 models we consider, identifying the characteristics of their underlying data generating processes. The column titled "Processes for r, g, & π" indicates the nature of the time-series models used to generate the interest rates, dividend growth rates, and equity premium. See Appendix 1 for details on how this set of models was chosen and a description of how the equity premium series is produced. The column titled "Downward Trend in Equity Premium Process," identifies whether the ex ante 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 basis point (bps) drop in the value of the ex ante equity premium. The column "Structural Break in Dividend Growth Process," indicates whether the model incorporates a gradual 100 basis point 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. Parsimonious Model: interest rates follow an AR(1), dividend growth rates follow a MA(1), the equity premium follows an AR(1). Deterministic π Model: interest rates follow an AR(1), dividend growth rates follow a MA(1), the equity premium follows a deterministic downward trend with a 50 bps structural break. Best BIC Model:† interest rates follow an ARMA(1,1), dividend growth rates follow a MA(1), the equity premium follows an AR(1). Second-Best BIC Model:† interest rates follow an ARMA(1,1), dividend growth rates follow a MA(1), the equity premium follows an ARMA(1,1). Further details about each model feature are provided in Appendix 1.

Model	Processes for r, g, & 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	Parsimonious Model	No	No	No	Yes
2	Parsimonious Model with π Trend	Yes (80 bps)	No	No	Yes
3	Parsimonious Model with π Break	No	Yes (50 bps)	No	Yes
4	Parsimonious Model with Dividend Growth Trend	No	No	Yes	Yes
5	Parsimonious Model with π Trend and Dividend Growth Trend	Yes (80 bps)	No	Yes	Yes
6	Parsimonious Model with π Break, π Trend, and Dividend Growth Trend	Yes (30 bps)	Yes (50 bps)	Yes	Yes
7	Best BIC Model† with π Break, π Trend, and Dividend Growth Trend	Yes (30 bps)	Yes (50 bps)	Yes	Yes
8	Second-Best BIC Model† with π Break, π Trend, and Dividend Growth Trend	Yes (30 bps)	Yes (50 bps)	Yes	Yes
9	Parsimonious Model with π Break and π Trend	Yes (30 bps)	Yes (50 bps)	No	Yes
10	Deterministic π Model with π Break and π Trend	Yes (30 bps)	Yes (50 bps)	No	Yes
11	Parsimonious Model with Constant Parameters π Break, π Trend, and Dividend Growth Trend	Yes (30 bps)	Yes (50 bps)	Yes	No
12	Parsimonious Model with Constant Parameters	No	No	No	No

† For 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.

Figure 1: Probability Distribution Functions of Simulated Ex Post Equity Premia, Dividend Yields, Sharpe Ratios, and Return Standard Deviations

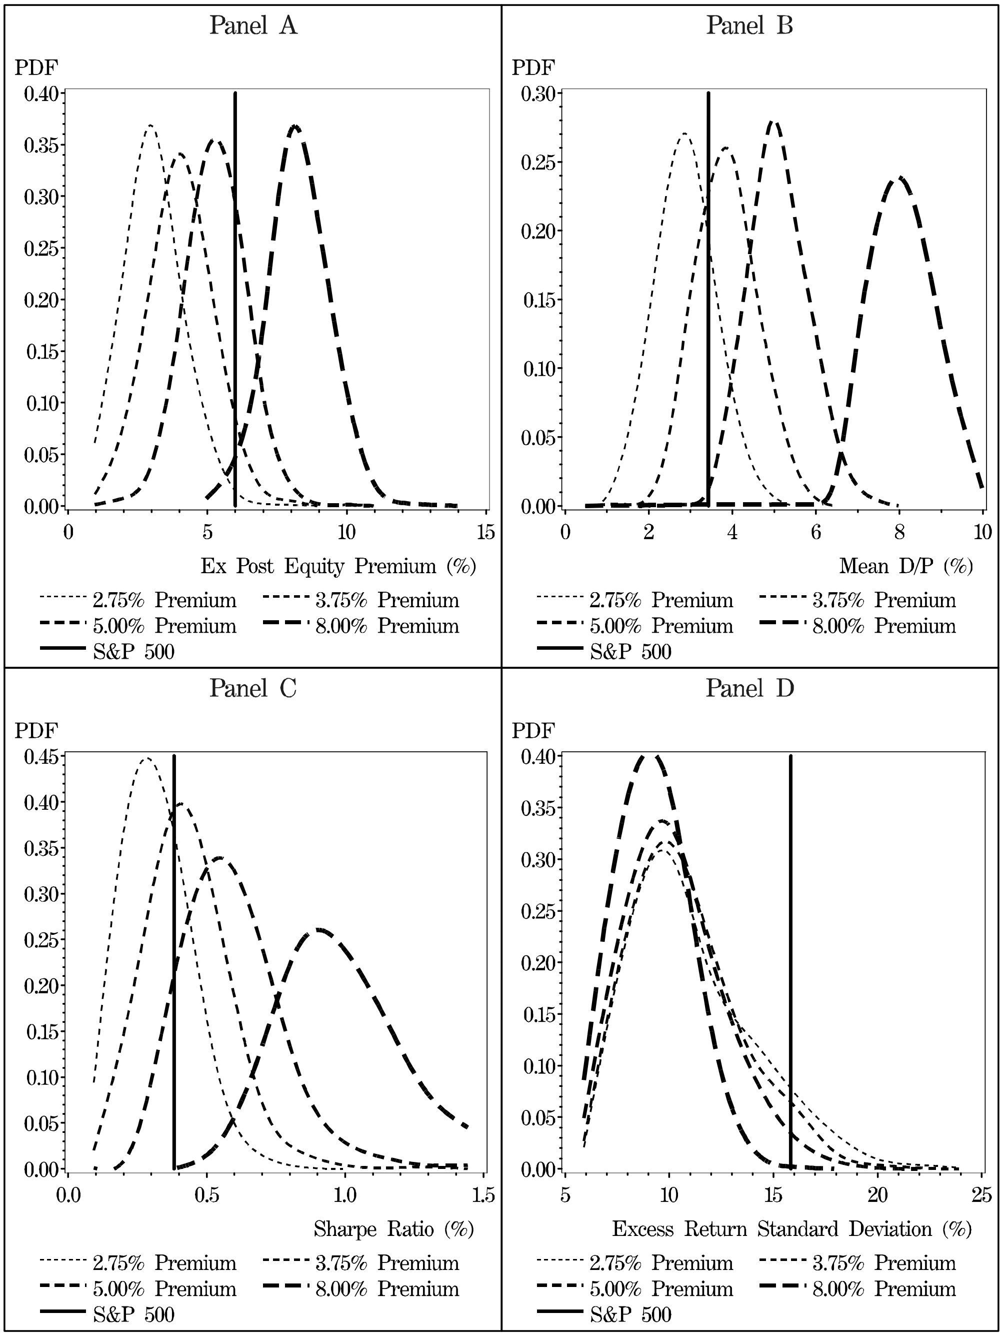

This figure contains probability distribution functions (PDFs) for various financial statistics generated in 2,000 simulated economies based on Model 1 from Table I. Each panel contains a PDF for each of four different assumed values of the ex ante equity premium: 2.75%, 3.75%, 5%, and 8%. Panel A shows the distribution of the ex post equity premium (mean return minus mean interest rate), Panel B shows the mean dividend yield distribution (dividend divided by price), Panel C shows the Sharpe ratio distribution (excess return divided by the standard deviation of the excess return), and Panel D shows the distribution of the standard deviation of excess returns. In each panel, a vertical line indicates the US 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 actual US results.

Figure 2: Joint and Individual Tests Statistics for Models 1 and 2

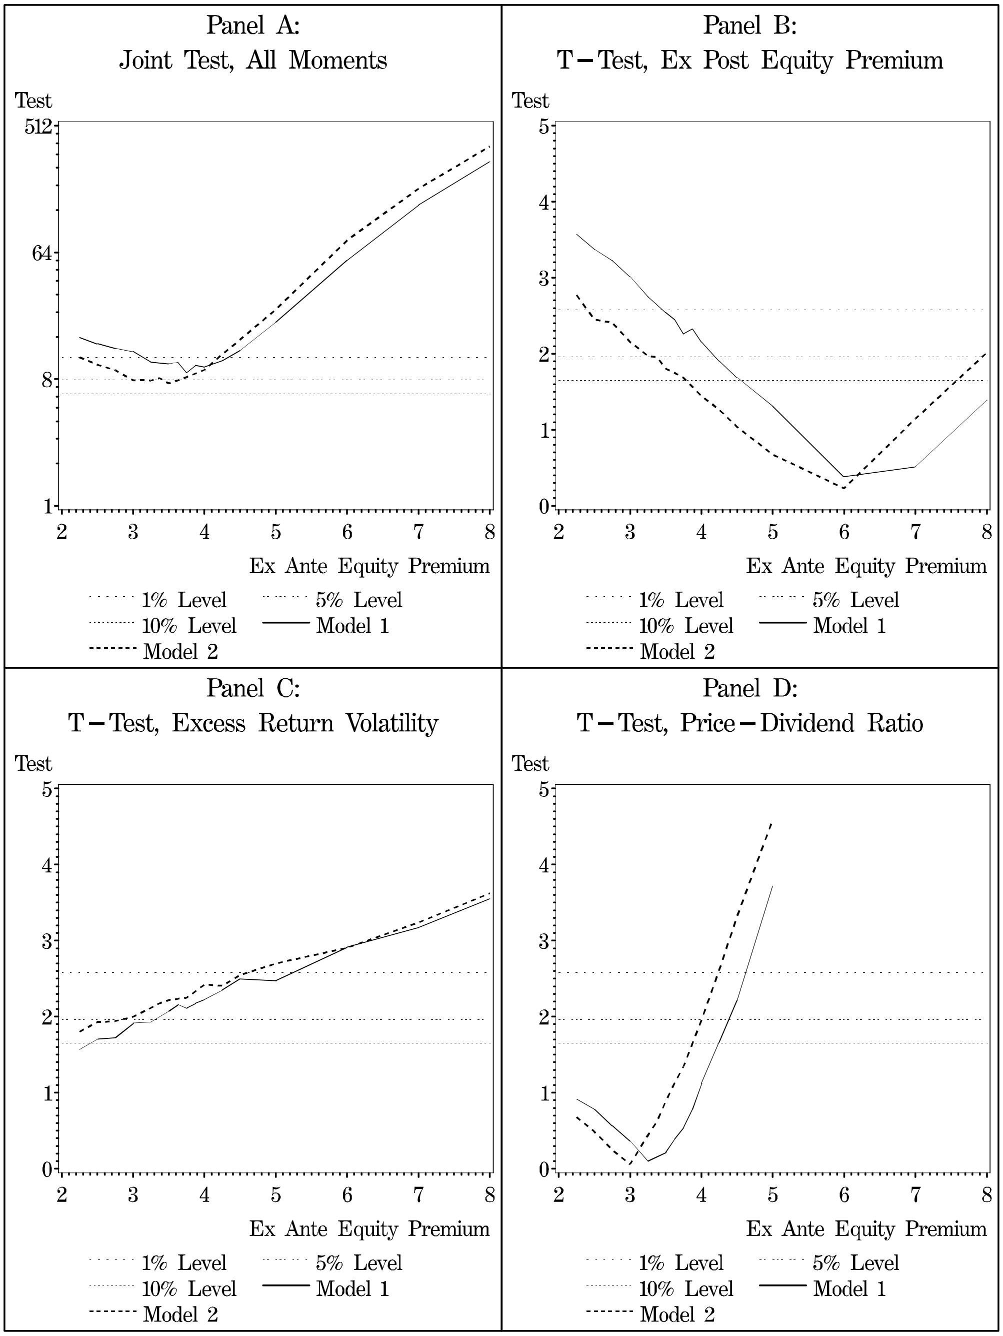

This figure contains plots of test statistics for Models 1 and 2. Panel A plots joint χ^2 tests based on a set of three variables (the ex post equity premium, the mean dividend yield, and the excess return volatility) for various ending values of the ex ante equity premium for each model. In Panel A the vertical axis is plotted on a log scale. The remaining panels contains t-test values corresponding to tests on the individual variables for each of the models: the ex post equity premium in Panel B, the excess return volatility in Panel C, and price-dividend ratio in Panel 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.

Figure 3: Joint and Individual Tests Statistics for Models 3 and 4

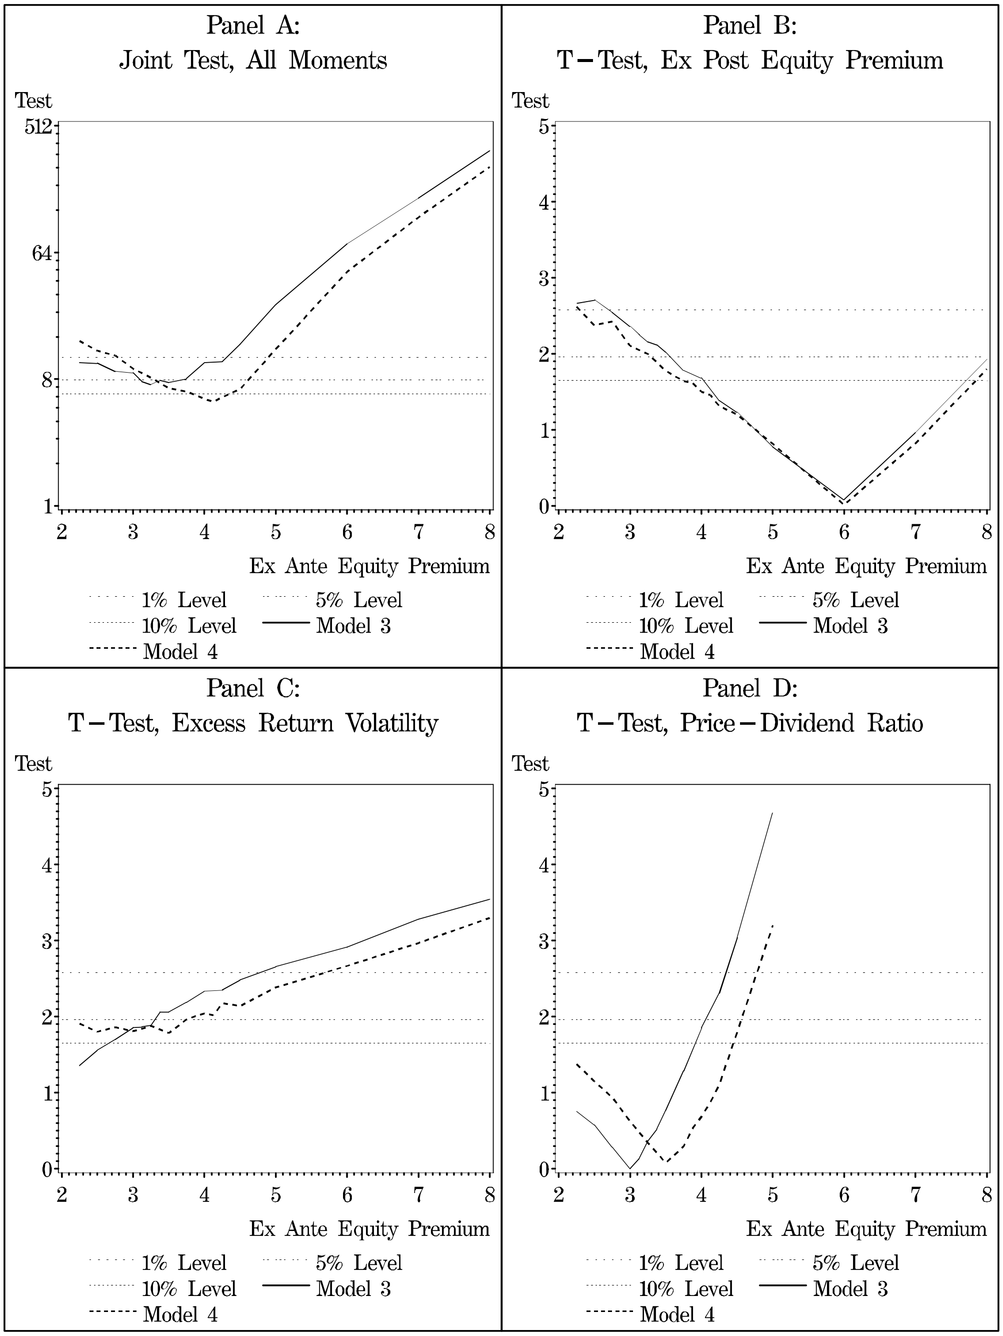

This figure contains plots of test statistics for Models 3 and 4. Panel A plots joint χ^2 tests based on a set of three variables (the ex post equity premium, the mean dividend yield, and the excess return volatility) for various ending values of the ex ante equity premium for each model. In Panel A the vertical axis is plotted on a log scale. The remaining panels contains t-test values corresponding to tests on the individual variables for each of the models: the ex post equity premium in Panel B, the excess return volatility in Panel C, and price-dividend ratio in Panel 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.

Figure 4: Joint and Individual Tests Statistics for Models 5 and 6

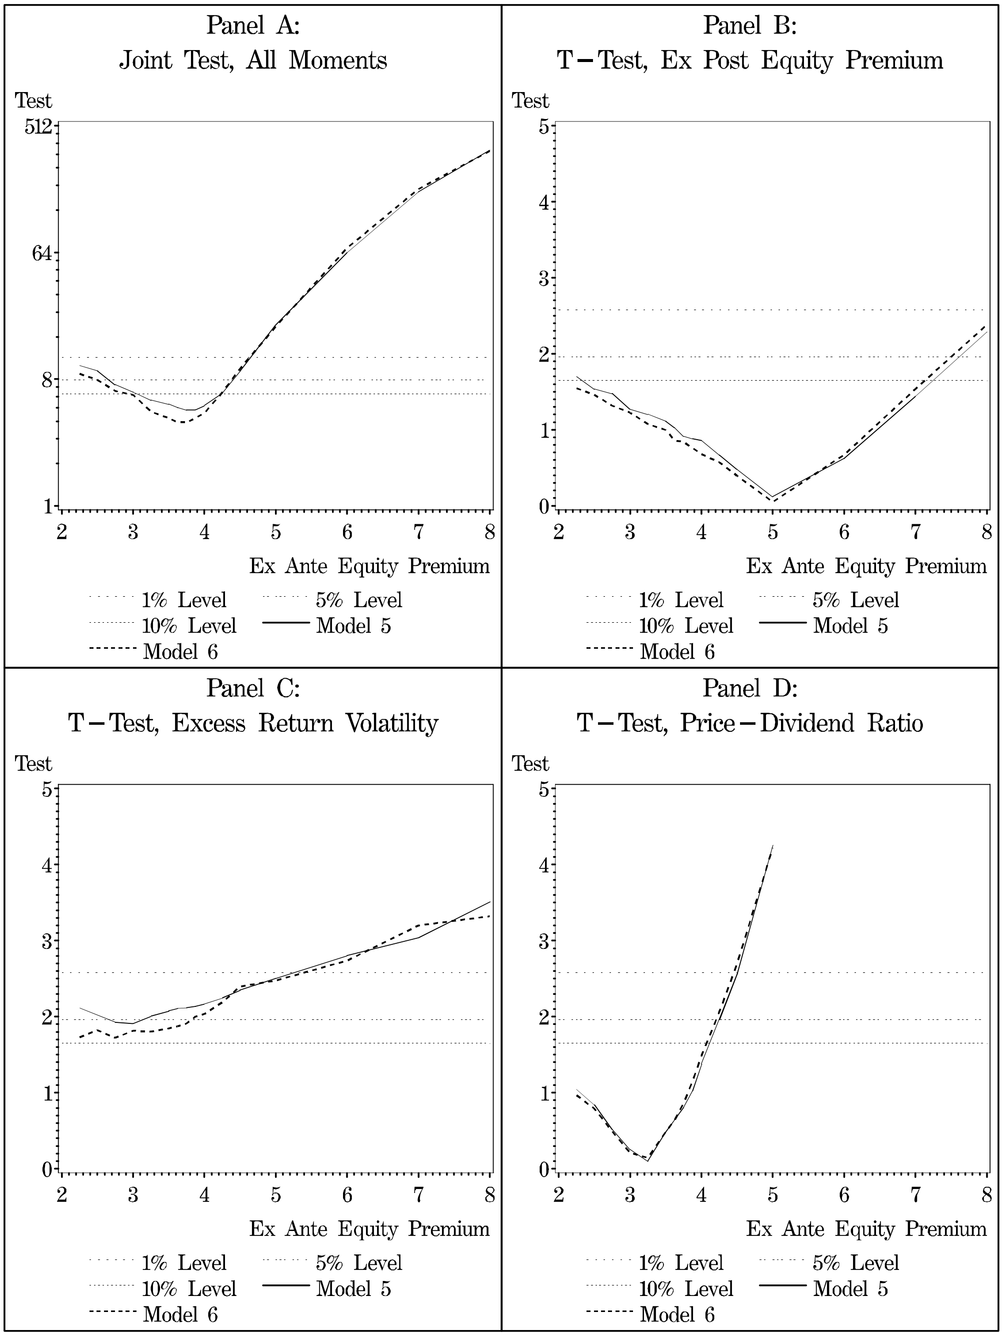

This figure contains plots of test statistics for Models 5 and 6. Panel A plots joint χ^2 tests based on a set of three variables (the ex post equity premium, the mean dividend yield, and the excess return volatility) for various ending values of the ex ante equity premium for each model. In Panel A the vertical axis is plotted on a log scale. The remaining panels contains t-test values corresponding to tests on the individual variables for each of the models: the ex post equity premium in Panel B, the excess return volatility in Panel C, and price-dividend ratio in Panel 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.

Figure 5: Joint and Individual Tests Statistics for Models 7 and 8

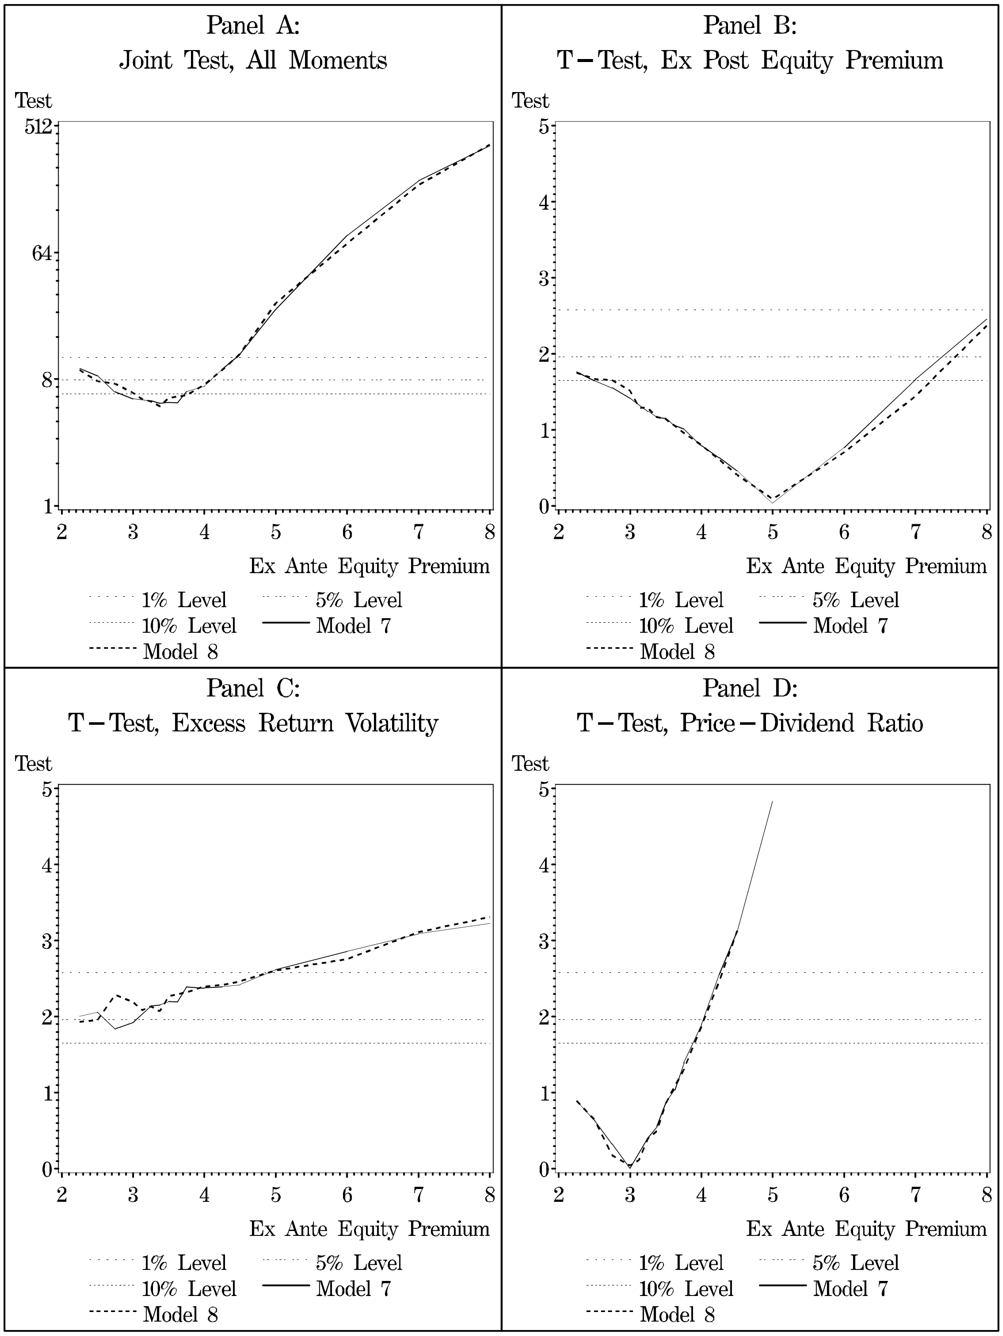

This figure contains plots of test statistics for Models 7 and 8. Panel A plots joint χ^2 tests based on a set of three variables (the ex post equity premium, the mean dividend yield, and the excess return volatility) for various ending values of the ex ante equity premium for each model. In Panel A the vertical axis is plotted on a log scale. The remaining panels contains t-test values corresponding to tests on the individual variables for each of the models: the ex post equity premium in Panel B, the excess return volatility in Panel C, and price-dividend ratio in Panel 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.

Figure 6: Joint and Individual Tests Statistics for Models 9 and 10

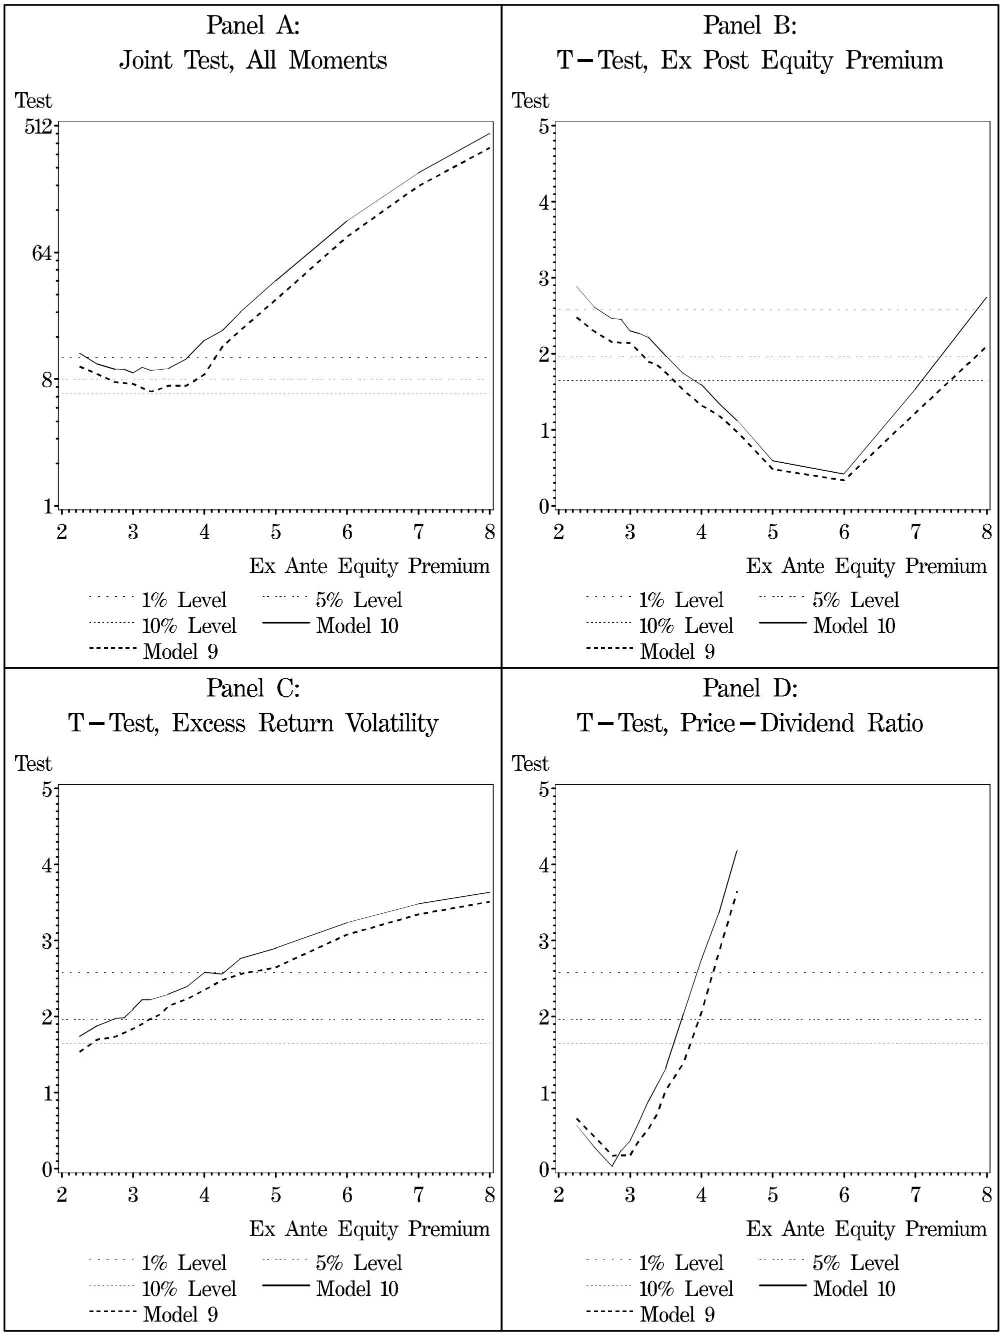

This figure contains plots of test statistics for Models 9 and 10. Panel A plots joint χ^2 tests based on a set of three variables (the ex post equity premium, the mean dividend yield, and the excess return volatility) for various ending values of the ex ante equity premium for each model. In Panel A the vertical axis is plotted on a log scale. The remaining panels contains t-test values corresponding to tests on the individual variables for each of the models: the ex post equity premium in Panel B, the excess return volatility in Panel C, and price-dividend ratio in Panel 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.

Figure 7: Parameter Estimation Certainty:
Joint and Individual Tests Statistics for Models 11 and 12

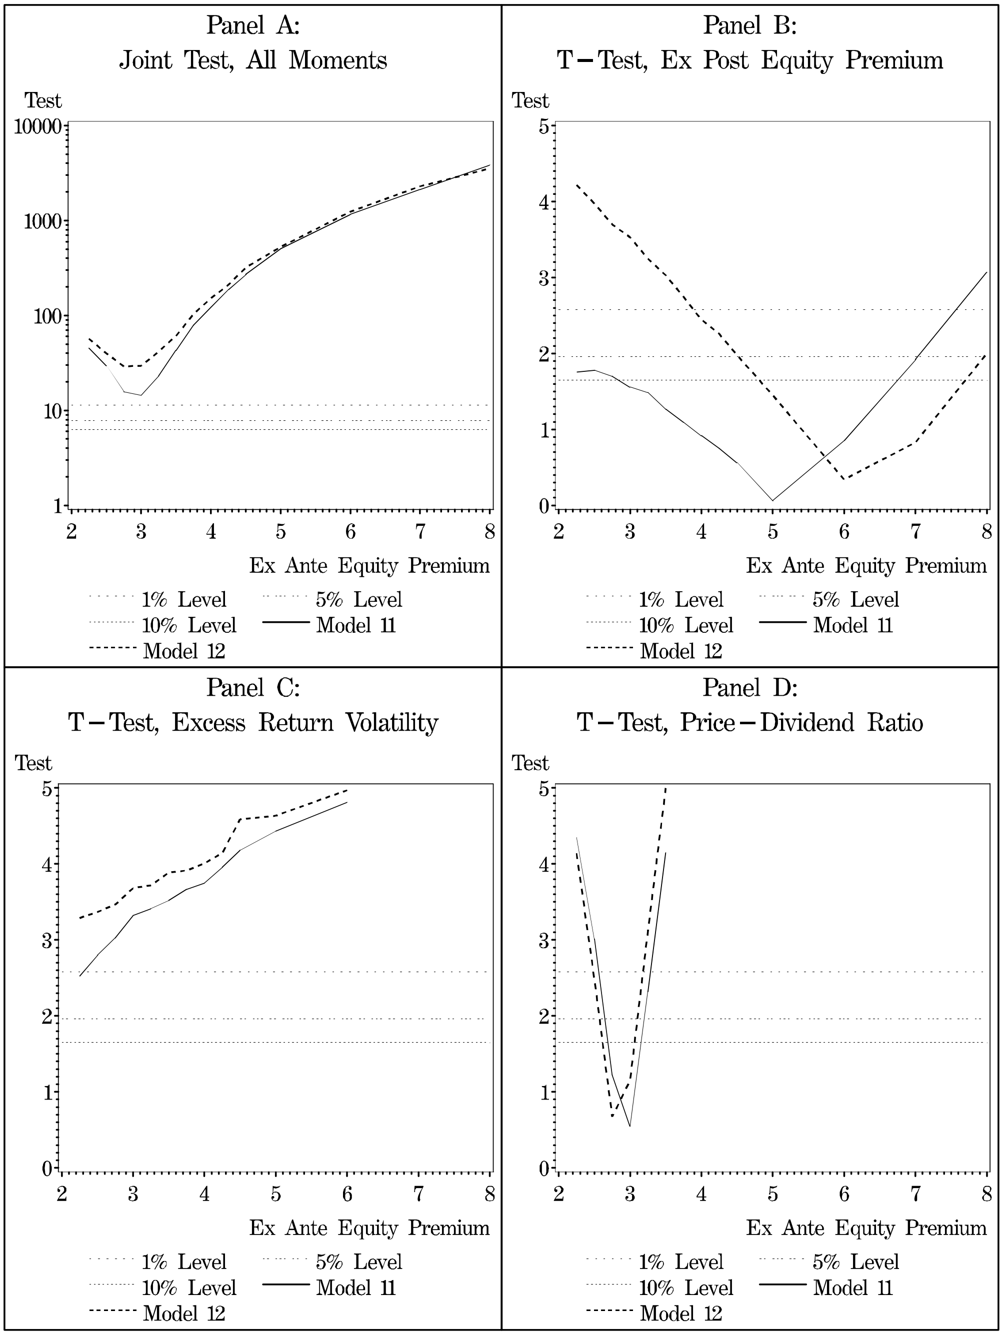

This figure contains plots of test statistics for Models 11 and 12. Panel A plots joint χ^2 tests based on a set of three variables (the ex post equity premium, the mean dividend yield, and the excess return volatility) for various ending values of the ex ante equity premium for each model. In Panel A the vertical axis is plotted on a log scale. The remaining panels contains t-test values corresponding to tests on the individual variables for each of the models: the ex post equity premium in Panel B, the excess return volatility in Panel C, and price-dividend ratio in Panel 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.

Figure 8: Investors' Model Uncertainty
Joint Tests Based on a Subset of Moments for Models 1-12

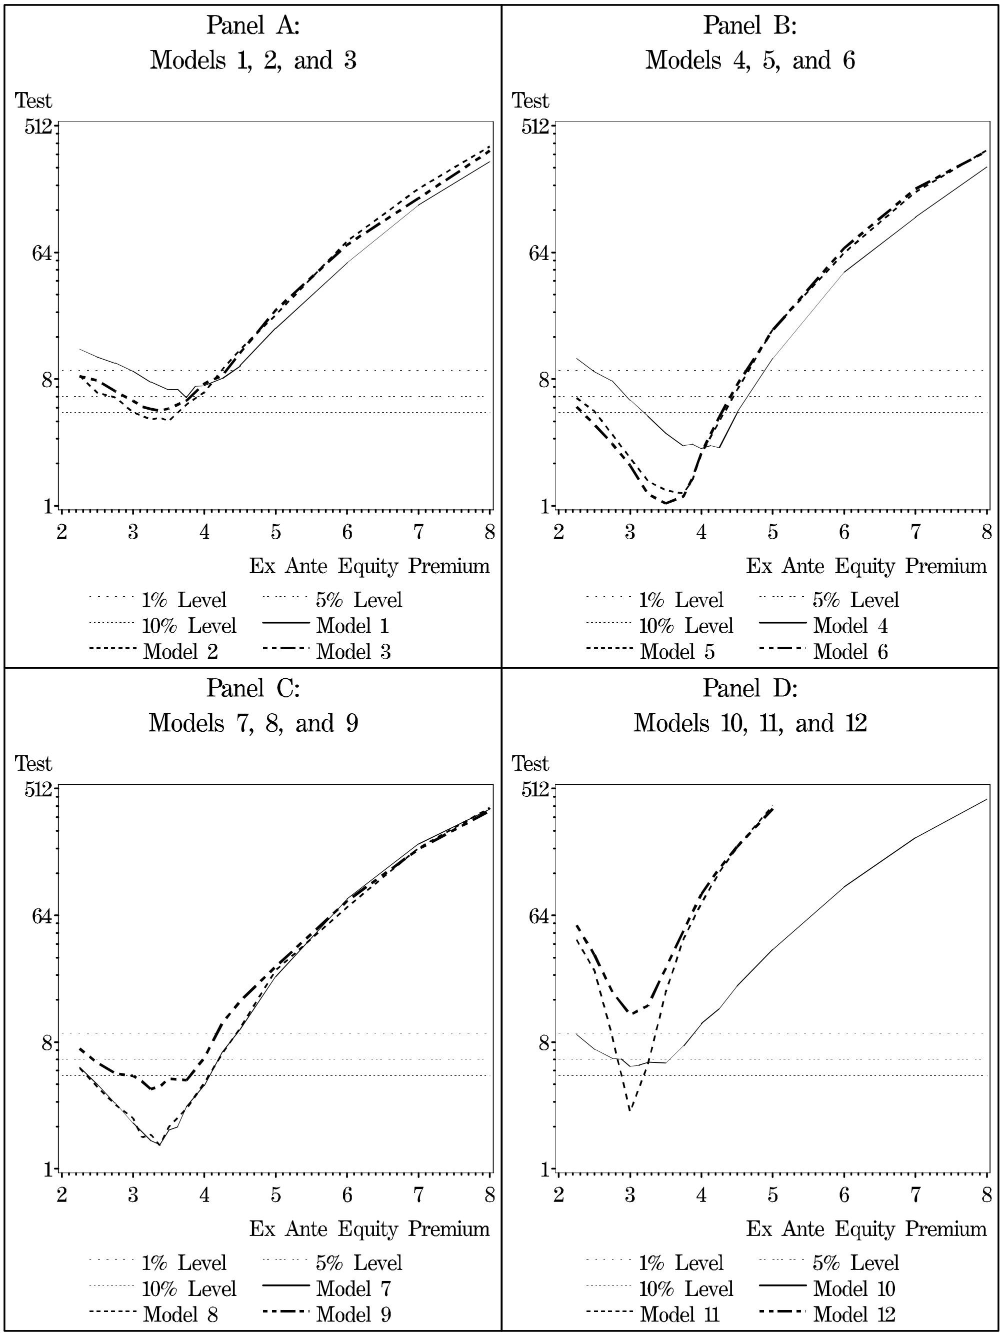

This figure contains plots of joint χ^2 tests based on a set of two variables, the ex post equity premium and the mean dividend yield, for various ending values of the ex ante equity premium for each model. Panel A presents the test statistics for Models 1, 2, and 3, Panel B presents the test statistics for Models 4, 5, and 6, Panel C presents the test statistics for Models 7, 8, and 9, and Panel D presents the test statistics for Models 10, 11, and 12. The vertical axis of each plot is on a log scale. 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.