Estimating and Testing Fundamental Stock Prices: Evidence from Simulated Economies

R. Glen Donaldson
Finning Ltd. Associate Professor of Finance
Faculty of Commerce and Business Administration
University of British Columbia
2053 Main Mall, Vancouver, BC, Canada, V6T 1Z2
Phone: 604-822-8344
Fax: 604-822-8521
E-Mail: glen@donaldson.commerce.ubc.ca.

and[1]

Mark Kamstra
Associate Professor of Economics
Department of Economics
Simon Fraser University
Burnaby BC, Canada V5A 1S6
Phone: 604-291-4514
Fax: 604-291-5944
Email: kamstra@sfu.ca.

Version: December 12, 2005

[1]We thank Lisa Kramer for helpful comments and the Social Sciences and Humanities Research Council of Canada for financial support.

A core tenet of financial economics states that, in a market populated by rational investors, the fundamental price of an asset equals the expected discounted present value of its future cashflows. This implies that in a rational and efficient market, and in the absence of price bubbles,[2] stock price movements are driven by forecasted changes in dividends and discount rates and not by the "irrational exuberance" of traders. Interest in the extent to which stock prices are fundamentally driven has been particularly high of late, with everyone from academics to the Chairman of the U.S. Federal Reserve Board offering opinions. Research into fundamental valuation is therefore particularly timely and, at least for the type of computer-intensive research we undertake in this chapter, is made possible by recent increases in computing power available to financial econometricians.

Several different procedures have been developed to estimate fundamental stock prices and to test whether market prices deviate in significant ways from the estimated fundamental price. Many studies have reported what appear to be important deviations from fundamentals, especially around market crashes such as those which occurred in 1929 and 1987. The resulting belief that financial markets may be "excessively volatile", and may potentially contain "price bubbles" that push market prices way from fundamental valuations, has contributed to the institution of policies such as trading halts in the face of large price moves. Whether such policies help or hurt financial markets depends in important ways on the extent to which market prices are driven by fundamental factors as opposed to "irrational exuberance", and the answer to this question crucially depends on the accuracy of the fundamental price estimates and tests used in the analysis.

To the best of our knowledge, no one has yet undertaken a thorough and systematic investigation of the accuracy of various fundamental price estimating and testing procedures.[3] Thus, while a fundamentals-estimation exercise might value a share of stock at

[2]The term "bubble" is formally defined below, but for now can be thought of as a situation in which some force other than a "fundamental factor", such as dividends and discount rates, drives stock prices.

[3]Bollerslev and Hodrick (1995) provide a survey of much of this literature, and themselves use simulation techniques to investigate tests of market efficiency. Bollerslev and Hodrick (1995) focus much of their simulation study on the constant discount rate case and consider only one ad hoc model for time-varying discount rates. This ad hoc model, like the constant discount rate case, allows an analytic solution for the

$X, we cannot be sure how confident to be in the estimate. For example, the estimate could be biased up or down and/or be very imprecise (i.e., have a large variance). With a small bias and low variance the estimate might be very accurate and therefore of considerable use. Conversely, a significantly biased estimate with very low precision might be wildly inaccurate and thus of little value. Establishing the properties of various fundamentals estimation and testing procedures would therefore be of significant benefit to academics, investors and policy makers who often employ fundamental price estimates in their work. The accuracy of fundamental price estimates can be obtained analytically under only very restrictive and special circumstances (such as assuming that dividend growth rates and interest rates will never change in the future, which is of course highly unrealistic). To circumvent this analytical intractability, we develop and test in this chapter a simulation-based method for calculating the properties of various fundamental price estimating and testing procedures. We are particularly interested in determining the statistical properties of fundamental price and return estimates commonly used in both industry and academia and in investigating the effects that estimation inaccuracies may have on the variety of volatility tests commonly employed in the literature.

Our strategy in this chapter is as follows: (a) use financial market data to estimate time-series models for dividend growth and discount rates, (b) use these models to simulate dividend growth and discount rate paths for a variety of possible economies that do not contain bubbles, (c) calculate fundamental prices for these bubble-free economies based on the simulated dividend growth and discount rates – which prices we call "market prices" since they are fair-value prices for the simulated market economies, (d) use various fundamental valuation models to estimate fundamental prices for each of the simulated economies, and (e) compare market prices versus fundamental prices and investigate statistical properties of common tests for excess volatility and bubbles in stock prices using

market price. We consider formal discount rate processes that admit serial dependence and hence rule out a general analytic solution. Allowing serial dependence in discount rates considerably complicates the analysis, but also considerably enriches the range of possible dynamics for prices and returns.

the data we simulated data under the (true) null hypothesis that there are no price bubbles. We apply this procedure to S&P 500 stock price data.

Results produced using our computer-aided techniques suggest that, while stock prices are indeed volatile, they are not more volatile than one would reasonably find in an economy driven by fundamentals, thereby suggesting that market prices are not "excessively volatile". We also find that traditional tests for excess volatility and bubbles over-reject a true null of no-bubbles in samples of the size traditionally employed in the literature. Indeed, most tests we investigate find overwhelming evidence of bubbles in data series we construct under the conditions that there are no bubbles. In other words, we demonstrate that traditional tests for price bubbles frequently find bubbles in data that do not in fact contain bubbles.

In Section I below we describe a variety of common fundamental pricing models and tests for excess volatility and bubbles. In Section II we present our Monte Carlo procedure for estimating fundamental prices/returns and investigating test performance. We discuss the results of our efforts in Section III. Section IV concludes.

I. Fundamental Pricing Methods and Tests

Define P_t^M as a stock's beginning-of-period-t market price, r_t as the rate used to discount payments received during period t, and $\mathcal{E}_t\left\{\cdot\right\}$ as the conditional expectation operator, with the conditioning information being the set of information available to investors at the beginning of period t. Investor rationality requires that the current market price of a stock, which will pay a dividend D_{t+1} at the beginning of period $t+1$ and then immediately sell for the ex-dividend market price P_{t+1}^M, satisfy Equation 1.

$$P_t^M = \mathcal{E}_t\left\{\frac{P_{t+1}^M + D_{t+1}}{1+r_t}\right\} \tag{1}$$

We can solve Equation 1 forward to period T (where $T > t$) and substitute realized dividends and discount rates in for their expected values to produce Equation 2.

$$P_t^X = \sum_{i=1}^{T-t} \left(\Pi_{k=1}^{i} \left[\frac{1}{1 + r_{t+k-1}} \right] \right) D_{t+i} + \left(\Pi_{k=1}^{T-t} \left[\frac{1}{1 + r_{t+k-1}} \right] \right) P_T^M \tag{2}$$

The superscript X on the left-hand-side price indicates that this is the ex post rational price; i.e., the price that an investor would have rationally paid for the stock had she known that the market price of the stock in period T was going to be P_T^M and that the stock was going to pay the sequence of dividends that it actually paid between periods t and T. The ex post rational price is not a fundamental price, nor is it a price that would ever be observed in the marketplace. However, for reasons explained below, it is nonetheless useful in analyzing price behavior.

I.A. Fundamental Pricing Models

To calculate the fundamental price of stock, we need to solve Equation 1 forward into infinite time under the transversality condition that the expected present value of the market price P_{t+i}^M falls to zero as i goes to infinity; i.e., there are no price "bubbles". This produces the familiar result that today's fundamental stock price, P_t^F, equals the expected present value of future dividends; i.e.,

$$P_t^F = \mathcal{E}_t \left\{ \sum_{i=1}^{\infty} \left(\Pi_{k=1}^{i} \left[\frac{1}{1 + r_{t+k-1}} \right] \right) D_{t+i} \right\}. \tag{3}$$

Note from Equation 3 that the fundamental price is a function only of dividends, discount rates and time, and not of the market price.

Defining the growth rate of dividends during period t as $g_t \equiv (D_{t+1} - D_t)/D_t$ allows the preceding equation to be rewritten as:

$$P_t^F = D_t \mathcal{E}_t \left\{ \sum_{i=1}^{\infty} \left(\Pi_{k=1}^{i} \left[\frac{1 + g_{t+k-1}}{1 + r_{t+k-1}} \right] \right) \right\}, \tag{4}$$

or, defining the discounted dividend growth rate as $y_t \equiv (1 + g_t)/(1 + r_t)$,

$$P_t^F = D_t \mathcal{E}_t \left\{ \sum_{i=1}^{\infty} \Pi_{k=1}^{i} y_{t+k-1} \right\}. \tag{5}$$

This can be rewritten as shown in Equation 6 below (note that y_t is *not* in the time t information set since it depends on D_{t+1} through g_t):

$$P_t^F = D_t \sum_{i=1}^{\infty} \mathcal{E}_t \left\{ \Pi_{k=1}^{i} y_{t+k-1} \right\}. \tag{6}$$

In practice, fundamental valuation calls for forecasting future discounted dividend growth rates (i.e., dividends and discount rates) to solve for the sum product in Equation 4 (or, equivalently, Equation 6). Since Equation 4 cannot be solved analytically in general, this usually requires that some restrictive assumptions be made concerning the time series processes driving dividends and discount rates. The simplest approach, introduced by Gordon (1962), is to assume that discount rates and dividend growth rates will be constant, at r and g respectively, for all future time so that Equation 4 reduces to Equation 7, in which the superscript G denotes the Gordon-model fundamental price estimate:

$$P_t^G = D_t \left[\frac{1 + g}{r - g} \right]. \tag{7}$$

The Gordon model is certainly convenient, but its extremely restrictive assumptions of constant r and g do not tend to produce the most accurate valuations possible. Several attempts have therefore been made to relax Gordon's original restrictions and yet still retain an analytical solution to Equation 4. The literature in this area is rather large and includes papers by Malkiel (1963), Fuller and Hsia (1984), Brooks and Helms (1990), Hurley and Johnson (1994, 1998) and Yao (1997), among others. Early literature in this area broke future time into several "chunks", with dividend growth and discount rates constant within each time-chunk, but different between chunks. For example, dividends might be forecasted to grow at a "high rate" for the first three years, and then grow at some "normal rate" for the rest of time, so that sum in Equation 4 would be broken into two parts, each part solved separately, and then added together (see, for example, Brooks and Helms (1990)).

Recent work on extending the Gordon model has largely focused on allowing for more complicated time-series behavior of dividends and discount rates, while retaining the ability to solve Equation 4 analytically. Two particularly good examples, found in Yao (1997), are the additive Markov model (Equation 1 of Yao (1997)) and the geometric Markov model (Equation 2 of Yao (1997)). These appear below as Equation 8 and Equation 9, respectively,

$$P_t^{ADD} = D_{t-1}/r + \left[1/r + (1/r)^2\right]\left(q^u - q^d\right)\Delta \tag{8}$$

$$P_t^{GEO} = D_{t-1}\left[\frac{1 + (q^u - q^d)\Delta^{\%}}{r - (q^u - q^d)\Delta^{\%}}\right] \tag{9}$$

in which q^u is the proportion of the time the dividend increases, q^d is the proportion of the time the dividend decreases, $\Delta = \sum_{t=2}^{T} |D_t - D_{t-1}|/(T-1)$ is the average absolute value of the level change in the dividend payment, and $\Delta^{\%} = \sum_{t=2}^{T} |(D_t - D_{t-1})/D_{t-1}|/(T-1)$ is the average absolute value of the percentage rate of change in the dividend payment.

Donaldson and Kamstra (1996, 2000) develop an alternative approach (hereafter referred to as the DK procedure) which does not require that Equation 4 be solved analytically.

The DK procedure instead uses Monte Carlo methods to solve Equation 4 numerically. The DK procedure eases the conditional constancy restrictions on dividend growth rates and discount rates of the Gordon procedure by modeling the discounted dividend growth rate – y in Equation 6 – explicitly as conditionally time-varying. The basic idea of the DK procedure is to estimate an econometric model for the time series behavior of y_{t+i} (the discounted dividend growth rate) in Equation 6 (which is equivalent to Equation 4), use this model and randomly-drawn innovations to simulate time-paths for the possible future evolution of y_{t+i}, then take the present value of the forecasted time-paths to find a fundamental price. This is done thousands (even millions) of times, with a different sequence of randomly-drawn innovations each time, so as to integrate out the expectation in Equation 6 and thus produce a numerical (as opposed to analytical) estimate of the fundamental price.

——— Exhibit 1 goes here ———-

The sequence of steps employed to produce a DK fundamental price estimate are described below and depicted in Exhibit 1.

Step A: Use in-sample data to specify and estimate an econometric model for conditionally time-varying y_t.

Step B: Use the estimated model from Step A and data up to period t to simulate, out-of-sample, possible realizations of y_{t+i} ; $i = 1, ..., I$, where I is chosen to be very large (infinity, were it practical),[4] for a cross section of J different possible economies, so that we follow the evolution of y across a panel of J simulated economies over I periods of time. We accomplish this by using the estimated model from Step A to make a conditional mean forecast of y, noted $\hat{y}_t$, then simulate a population of J possible shocks around the mean

[4]Note that as I increases, the product $\Pi_{i=1}^{I} y_{t+i}$ converges to zero since y is less than unity in steady state. We found that in practice $I = 400$ was easily sufficient to have our simulations converge to the point where increasing I had no impact on the fundamental price. In other words, a dividend received 400 years in the future has essentially no impact on today's price and thus 400 years in the future is equivalent to infinity for practical purposes in our simulations.

and add these shocks to $\hat{y}_t$ to produce y_t^j ; $j = 1, ..., J$ for the cross section of j economies at time t. We then repeat this procedure through time for each of the J economies, conditioning on y_t^j, to form y_{t+i}^j $i = 1, ..., I$. We therefore produce an out-of-sample panel of values of y for J economies stretching out I periods, all of which have been simulated based on the in-sample data.

Step C: Calculate the fundamental DK price based on this panel of simulated y using Equation 6 as follows:

$$P_t^{DK} = D_t \sum_{j=1}^{J} \left(\sum_{i=1}^{I} \Pi_{k=1}^{i} y_{t+k}^j \right) / J. \tag{10}$$

Note that the preceding equation can be rearranged by bringing the dividend level D_t inside the parentheses so that

$$P_t^{DK} = \sum_{j=1}^{J} \left(D_t \sum_{i=1}^{I} \Pi_{k=1}^{i} y_{t+k}^j \right) / J.$$

It is interesting to note here that the parenthesized term in the preceding equation – i.e., $D_t \sum_{i=1}^{I} \Pi_{k=1}^{i} y_{t+k}^j$ – is for economy j the ex post rational price of the stock from Equation 2, with the terminal period pushed out to infinity as I goes to infinity. In other words, if an investor stood at the end of time and looked back over history to see that the stock had paid a stream of dividends that had grown at rate g_{t+i}^j and which were discounted at rate r_{t+i}^j such that the discounted dividend growth rate had been y_{t+i}^j in this j^{th} economy, the investor would rationally feel that she should have been willing to pay $P_t^{X,j} \equiv D_t \sum_{i=1}^{I} \Pi_{k=1}^{i} y_{t+k}^j$ to have purchased the stock back in time t, where $P_t^{X,j}$ is commonly referred to by Shiller (1981) and others as the "ex post rational price" (although when applied to data it is truncated at the end of the data period, as shown in Equation 2 above).

The representation of a single fan of the DK simulation as an ex post rational price is interesting because it highlights the relationship between the ex post rational price, the Gordon price and the DK technique, and therefore provides some insight into the likely properties of the price estimates produced by various methods. For example, it is clearly seen that all Gordon-based methods are restricted versions of the DK technique. The basic Gordon model, for example, imposes conditionally constant dividend growth rates and discount rates such that each of the j economies in the simulation are identical and based on a constant y. If these restrictions are invalid – and there is good reason to believe that they are indeed invalid – then the Gordon estimator will be be biased and inefficient relative to DK.

The ex post rational price calculations performed by Shiller (1981) and others are also problematic because they estimate a fundamental price based only on one realization of dividends and discount rates: the realization actually observed in the true market. In other words, Shiller's method stands at the end of time and asks what an investor with perfect foresight would have paid back at date t for a share of stock, had she known the dividends and discount rates that were to occur in the future.[5] Conversely, a real-life investor looks into an uncertain future when making purchasing decisions and therefore considers a universe of possible economies that might unfold when valuing a share of stock – this is the DK method. The ex post rational price, based on only one time-path, will therefore provide an imprecise picture of the thousands of possible economies, and their associated dividend time-paths, that were considered when forming the market price at which real-life traders buy and sell stock. The extent to which this imprecision might affect our view on important asset-pricing questions, such as whether market prices are excessively volatile or not, is studied below.

[5] Of course, if the investor had perfect foresight, and therefore faced no uncertainty, she would not require an equity risk premium so the discount rate she would use would be lower than the rate r used in the ex-post price calculation. This subtlety is typically ignored in the literature that uses ex-post price calculations. For consistency we will also follow the literature here.

I.B. Some Tests Using Fundamental Prices

One of the key features of analytically-based fundamental pricing procedures is that the fundamental prices they produce behave differently than observed market prices in important ways. For example, the time path of such fundamental prices is typically much smother than that of observed market prices, which has led many researchers to conclude that market prices are excessively volatile. Indeed, there is an entire literature devoted to the study of excess volatility in financial markets (reviewed by Camerer (1989) and Cochrane (1992)). Conversely, Donaldson and Kamstra (1996, 2000) find that the DK procedure produces fundamental stock prices that behave substantially the same as market prices in terms of return volatility. Of course, the evidence presented in all of these papers is based on the comparison of one estimate of the fundamental price (the estimate produced by the model in question) with one realization of the market price (the price series we see in the true market data). Given that the results come from only one price series, there is no way to be truly sure how accurate (biased, precise, etc.) the fundamentals estimate might be in general. The use of computer-aided financial econometric techniques in an investigation of the accuracy of various fundamentals-estimation procedures is therefore one important objective of this chapter. Another objective of this chapter is to investigate the properties of various tests applied to fundamental and actual stock price data, in particular tests for excess volatility and price bubbles. Tests we investigate in this chapter include the following standard tests for price bubbles. First, Camerer (1989) observes that market prices would boom upward and crash downward with bubbles expanding and collapsing, which would lead market prices to be excessively volatile relative to fundamental prices if the market price contained bubbles. This suggests a test with a no-bubbles null hypothesis that the variance of the percentage rate of change in market prices and the percentage rate of change in fundamental prices are equal, and an alternative hypothesis that the market price contains bubbles and thus that the variance in the percentage rate of change in market prices is greater than the percentage rate of change in fundamental prices.

Second, a direct test of bubbles asks whether the market and fundamental prices are cointegrated (e.g., Campbell and Shiller (1987)). Bubbles in the market price will lead to a non-stationary difference between the market price and the fundamental price estimate; a standard unit root test on this difference therefore tests the null of cointegration. See Dickey and Fuller (1981) and Said and Dickey (1984) for a description of unit root tests.

Third, Mankiw, Romer and Shapiro (1985) (MRS) develop two tests for bubbles. They exploit a decomposition of the difference between the ex post rational price, P_t^X (calculated by inserting realized dividends and discount rates into the present value equation as in Shiller (1981), as seen in Equation 2 above), and the fundamental price estimate P_t^F, relative to the observed market price P_t^M:

$$\left(P_t^X - P_t^F\right) = \left(P_t^X - P_t^M\right) + \left(P_t^M - P_t^F\right).$$

Since the total volatility of the two sides of the preceding equation are equal by definition, the volatility of each individual component on the right hand side should be less than or equal to the volatility of the left hand side. The test which compares the volatility of the first term on the right hand side with the volatility of the left hand side will be labeled the MRS1 test, while the test which compares the volatility of the second term on the right hand side with the volatility of the left hand side will be labeled the MRS2 test. In both cases the null hypothesis of no bubbles (i.e., no excess volatility) is that the term on the left hand side is more volatile than the right hand side term to which it is being compared.

In considering these tests it is worth asking whether any of them have proper size, or whether they over-reject as some researchers argue. Note that these tests involve a joint null hypothesis that the market price does not contain a bubble and that the fundamental price *estimate*, produce by some fundamental pricing model, shares similar properties with the market's true unobserved fundamental price. A finding of systematic over-rejection across simulated market economies (where the null of no bubbles is imposed) would therefore indicate that the model used to produce fundamental price estimates is misspecified.

In this chapter we employ Monte Carlo procedures to investigate the performance of various fundamentals-estimation methodologies and of simulated market prices, and also of various procedures commonly used to test for excess volatility and bubbles in stock prices. Although our particular application of Monte Carlo methods is new, there is already a significant literature that uses Monte Carlo methods in asset pricing applications. Indeed, there is body of work that (directly or indirectly) simulates stock prices and dividends, under various assumptions, to investigate price and dividend behavior (e.g., Scott (1985), Kleidon (1986), West (1988a,b), Campbell (1991), Mankiw, Romer and Shapiro (1991), Hodrick (1992), Timmermann (1993,1995), and Campbell and Shiller (1998)). However, these studies typically impose restrictions on the dividend and discount rate processes so as to obtain fundamental prices from some variant of the Gordon (1962) model discussed above and/or some log-linear approximating framework.

Rather than impose approximations to solve Equation 4 analytically, we will instead simulate the dividend growth and discount rate processes directly, and evaluate the expectation through Monte Carlo integration techniques. This approach is computationally burdensome since it requires us to perform a Monte Carlo simulation of a Monte Carlo simulation, but it is the only way to evaluate Equation 4 without approximation error.[6] We also take care to calibrate our models to the time series properties of the data. Dividend growth, for instance, is strongly autocorrelated in the S&P500 stock market data, in contrast to the assumption of a log random walk for dividends often imposed in this literature.

[6] There is still Monte Carlo simulation error, but that is random, unlike most types of approximation error, and it can also be measured explicitly. This simulation error is analogous to simple cases such as the simulation error associated with Monte Carlo experiments on the size of a test statistic.

II. Simulating Fundamental Stock Prices

Consider again the pricing relationship of Equation 4 above, rewritten below for convenience, in which P is the price (for notational convenience we drop the superscript "F"), D the dividend, g the dividend growth rate, and r the discount rate:

$$P_t = D_t \mathcal{E}_t \left\{ \sum_{i=1}^{\infty} \left(\Pi_{k=1}^{i} \left[\frac{1 + g_{t+k-1}}{1 + r_{t+k-1}} \right] \right) \right\}.$$

We now execute our strategy of (a) calibrating models to market dividends and interest rates, (b) simulating bubble-free economies using these calibrated models, (c) calculating fundamental prices for these bubble-free economies, (d) forming fundamental price *estimates*, including the Gordon models and the DK model, and (e) evaluating these fundamental price estimates and tests for bubbles.

II.A. Dividends and Discount Rates

The first step is to estimate time series models for dividend growth and interest rates so that the Monte Carlo simulations generate dividends and discount rates that match real-world dividends and discount rates. This will allow us to generate non-bubble prices and returns and compare them to real world prices and returns (e.g, a finding that the real world prices and returns look significantly different than the simulated (non-bubble) prices and returns would provide evidence of non-fundamental movements in real world prices and returns; possibly bubbles). This will also allow us to evaluate conventional fundamental price estimation methods, including variants of the Gordon Growth model, and conventional tests for bubbles, such as the MRS tests, to explore the properties of the estimators and test statistics when there are no bubbles in the price.

Our dividend process is calibrated to the S&P 500 stock index annual dividend data, 1952-1998, collected as described in Shiller (1989). The discount rate is defined to be the risk free interest rate plus a constant equity premium, where the risk free rate is the interest rate on a one-year U.S. T-bill as constructed by the U.S. Federal reserve, 1952-1998.[7] A constant equity premium of 5.77% is added to the risk free interest rate to produce a discount rate consistent with the stock price data.[8]

S&P 500 dividends grew at an average rate of 5.5% per year, over 1952-1998, with a standard deviation of 3.7%. As dividend growth rates have a minimum value of -100% and no theoretical maximum, a natural choice for their distribution is the log normal in the S&P500 data. The logarithm of 1 plus the dividend growth rate has mean 0.0531 and standard deviation 0.035 over 1952-1998. We estimated simple autoregressive moving average (ARMA) time series models for the logarithm of 1 plus the dividend growth rate and found the best model by the Bayesian Information Criterion to be a moving average model of order 1 (MA(1)) with the MA(1) coefficient equal to 0.60. Standard tests for normality of this error term do not reject the null of normality,[9] and standard tests for autocorrelation and autoregressive conditional heteroskedasticity (ARCH) fail to reject the null of homoskedasticity and no serial correlation[10].[11]

As for interest rates, since economic theory admits a wide range of possible interest rate processes there are a variety of models possible, from constant to autoregressive and highly non-linear heteroskedastic forms. The autoregressive model of order 1 (AR(1)) of the logarithm of interest rates, as described in Hull (1993) p.408, will be used here as it fits our data well and restricts nominal rates to be positive. Standard specification tests for

[7]We choose this measure because, as Cochrane (1992) notes, "there is a long tradition in the volatility test and investment or capital-budgeting literature that measures time-varying discount rates from interest rates plus risk premiums that are constant over time".

[8]Finance theory requires that $\mathcal{E}\left\{(1+R)/(1+i+\pi)\right\} = 1$, where R is the log of one plus the stock return, i the riskfree rate and π the equity premium (i.e., $1/[1+i+\pi]$ is the "pricing kernel"). See, for example, Campbell, Lo and MacKinlay (1997). Selection of the equity premium that produces the requisite pricing kernel was accomplished with a grid search, and this leads to a risk premium of 5.77% in the annual S&P 500 data over 1952-1998.

[9]These tests include the Shapiro-Wilk, Kolmogorov-Smirnov, Cramer-von Mises, and Anderson-Darling tests. See SAS Procedures Guide (1999).

[10]See Engle (1982) for the seminal treatment of ARCH effects.

[11]These are based on autocorrelations of the residuals and residuals squared.

normality, autocorrelation and ARCH on the error term from an AR(1) model of the logarithm of interest rates do not reject the null of no misspecification. The 1-year T-bill rates have mean 0.059 and standard deviation 0.03 over 1952-1998. The AR(1) coefficient estimate in the regression of log interest rates on lagged log interest rates equals 0.83. Finally, the error terms from the MA(1) model of log dividend growth rates and log interest rates are correlated, with a correlation coefficient of 0.21.

Properties of fundamental prices and returns produced by Equation 4 hinge delicately on the modeling of the dynamics of the dividend growth and interest rate processes. For instance, fundamental prices will equal a constant times the dividend level and fundamental returns will be very smooth over time if dividend growth and interest rates are equal to constants plus independent innovations. However, modeling these data series to capture the serial dependence of dividend growth rates and interest rates observed in the data, as we have done, will typically lead to time-varying price-dividend ratios and variable returns of the sort we see in the S&P500 stock market data.

II.B. The Monte Carlo Experiment

We now detail the Monte Carlo experiment by which the price, P, is arrived at given conditioning information on the dividend level, D, dividend growth rate, g, and interest rate, r. That is, we detail for the e^{th} economy (where $e = 1, ..., E$) the formation of the price P_t^e given D_t^e, g_t^e, and r_t^e. P_t^e is the market price that Equation 4 states would obtain in economy e if the stock market in economy e is rational, efficient, and bubble-free.

In terms of timing and information, recall that P_t^e is the stock's beginning-of-period-t market price based on fundamental factors, r_t^e is the rate used to discount payments received during period t and is known at the beginning of period t, D_t^e is paid at the beginning of period t, $g_t^e \equiv (D_{t+1}^e - D_t^e)/D_t^e$ and is not known at the beginning of period t since it depends on D_{t+1}^e, and $\mathcal{E}_t\{\cdot\}$ is the conditional expectation operator, with the conditioning information being the set of information available to investors at the beginning of period t. Finally, recall from Equation 4 that investor rationality requires

$$P_t^e = D_t^e \mathcal{E}_t \left\{ \sum_{i=1}^{\infty} \left(\Pi_{k=1}^{i} \left[\frac{1 + g_{t+k-1}^e}{1 + r_{t+k-1}^e} \right] \right) \right\}.$$

Based on the preceding equation, we generate prices by generating a multitude of possible streams of dividends and discount rates, present-value discounting the dividends with the interest rates, and averaging the results; i.e., by conducting a Monte Carlo integration. Hence we produce $P_t^e,\ e = 1, ..., E$ utilizing only dividend growth rates and discount rates. The exact procedure is described below and summarized in Exhibit 2.

———— Exhibit 2 goes here ————-

II.B.1. Simulating Stock Prices for a Rational, Efficient, Bubble-Free Market

Step 1: When forming P_t^e, the most recent fundamental information available to a market trader would be g_{t-1}^e, D_t^e, and r_t^e. The quantities g_{t-1}^e, D_t^e, and r_t^e must therefore be generated directly in our Monte Carlo experiment, whereas P_t^e must be calculated based on these g, r and D since this is how prices would be determined by a rational market participant in a bubble-free economy. In the steps below the risk free T-bill rate is indicated as r and the discount rate (i.e., the risk free interest rate plus the risk premium of 5.77%) as r_*. The objective of Steps 1(a)-(c) is to produce dividend growth rates and interest rates that replicate the real world dividend growth and discount rate data. That is, the simulated dividend growth rates and interest rates must have the same mean, variance, correlation structure and autocorrelation structure as the real world dividend growth rates and interest rates.

Step 1(a): Note that since the logarithm of one plus the dividend growth rate is modeled as a MA(1) process, $log(1+g^e_t)$ is a function of only innovations, labeled ϵ^e_g. Note also that since the logarithm of the interest rate is modeled as an AR(1) process, $log(r^e_t)$ is a function of $log(r^e_{t-1})$ and an innovation labeled ϵ^e_r. Set the initial dividend, D^e_1, equal to the S&P500's dividend value for 1951 (observed at the end of 1951), and the lagged innovation of the logarithm of the dividend growth rates $\epsilon^e_{g,0}$ to 0. To match the real-world interest rate data, set $log(r^e_0) = -3.05$ (the mean value of log interest rates required to produce interest rates matching the mean and variance of observed T-bill rates). Set the standard deviation of the innovation to the log interest rate process to 0.242, and the standard deviation of the innovation to the log dividend growth rate process to 0.0305. Then generate two independent standard normal random numbers, $\epsilon^e_{1,1}$ and $\epsilon^e_{2,1}$, and form two correlated random variables, $\epsilon^e_{r,1} = 0.242(0.21\epsilon^e_{1,1} + (1-.21^2)^{.5}\epsilon^e_{2,1})$ and $\epsilon^e_{g,1} = 0.0305\epsilon^e_{1,1}$. These are the simulated innovations to the interest rate and dividend growth rate processes, formed to have standard deviations of 0.242 and 0.0305 respectively to match the data, and to be correlated with correlation coefficient 0.21 as we find in the S&P 500 and T-bill data. Next, form $log(1+g^e_1) = 0.0531 + 0.60\epsilon^e_{g,0} + \epsilon^e_{g,1}$ and $log(r^e_1) = -0.18 + 0.94log(r^e_0) + \epsilon^e_{r,1}$.[12] Also form $D^e_2 = D^e_1(1+g^e_1)$.

Step 1(b): Produce two correlated normal random variables, $\epsilon^e_{r,2}$ and $\epsilon^e_{g,2}$ as in Step 1(a) above, and conditioning on $\epsilon^e_{g,1}$ and $log(r^e_1)$ from Step 1(a) produce $log(1+g^e_2) = 0.0531 + 0.60\epsilon^e_{g,1} + \epsilon^e_{g,2}$, $log(r^e_2) = -0.18 + 0.94log(r^e_1) + \epsilon^e_{r,2}$ and $D^e_3 = D^e_2(1+g^e_2)$.

Step 1(c): Repeat Step 1(b) to form $log(1+g^e_t)$, $log(r^e_t)$ and D^e_t for $t = 3, 4, 5, ..., T$ and for each economy $e = 1, 2, 3, ..., E$. Then calculate the dividend growth rate g^e_t and the discount rate $r^e_{*,t} = r^e_t + 0.0577$.

Step 2: For each time period $t = 1, 2, 3, ..., T$ and economy $e = 1, 2, 3, ..., E$ we must

[12] Notice that the AR(1) parameter for the log interest rate process is estimated to be 0.83 but we have set it to 0.94 in the simulations. It is well known that the coefficient estimate in an AR(1) OLS regression is biased downwards; see for instance Kennedy (1992) p.147. Monte Carlo experiments were employed to determine the appropriate correction for our data, as in Orcutt and Winokur (1969), and this led to the setting of 0.94. The intercept term had to be adjusted as well to reflect this new setting.

calculate present value prices, P_t^e. 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 e. That is, we must produce a cross-section of dividends and interest rates that might be observed in periods $t, t+1, t+2, ...$ given what is known at period $t-1$ and use these to solve the expectation of Equation 4. The counter j below indexes the cross-section of future economies that could possibly evolve from the current state of the economy.

Step 2(a): Set $\epsilon_{g,t-1}^{j,e} = \epsilon_{g,t-1}^{e}$ and $log(r_{t-1}^{j,e}) = log(r_{t-1}^{e})$ for $j = 1, 2, 3, ..., J$. Generate two independent standard normal random numbers, $\epsilon_{1,t}^{j,e}$ and $\epsilon_{2,t}^{j,e}$ and form two correlated random variables $\epsilon_{r,t}^{j,e} = 0.242(0.21\epsilon_{1,t}^{j,e} + (1-.21^2)^{.5}\epsilon_{2,t}^{j,e})$ and $\epsilon_{g,t}^{j,e} = 0.0305\epsilon_{1,t}^{j,e}$ for $j = 1, 2, 3, ..., J$.[13] These are the simulated innovations to the interest rate and dividend growth rate processes, respectively. Form $log(1+g_t^{j,e}) = 0.0531 + 0.60\epsilon_{g,t-1}^{j,e} + \epsilon_{g,t}^{j,e}$ and $log(r_t^{j,e}) = -0.18 + 0.94log(r_{t-1}^{j,e}) + \epsilon_{r,t}^{j,e}$.

Step 2(b): Produce two correlated normal random variables $\epsilon_{r,t+1}^{j,e}$ and $\epsilon_{g,t+1}^{j,e}$ as in Step 2(a) above, and conditioning on $\epsilon_{g,t}^{j,e}$ and $log(r_t^{j,e})$ from Step 2(a) produce $log(1+g_{t+1}^{j,e}) = 0.0531 + 0.60\epsilon_{g,t}^{j,e} + \epsilon_{g,t+1}^{j,e}$, and $log(r_{t+1}^{j,e}) = -0.18 + 0.94log(r_t^{j,e}) + \epsilon_{r,t+1}^{j,e}$ for $j = 1, 2, 3, ..., J$.

Step 2(c): Repeat Step 2(b) to form $log(1+g_{t+i}^{j,e})$ and $log(r_{t+i}^{j,e})$ for $i = 2, 3, 4, ..., I$, $j = 1, 2, 3, ..., J$, and economies $e = 1, 2, 3, ..., E$. Solve for the dividend growth rate $g_{t+i}^{j,e}$, the dividends $D_{t+i}^{j,e}$, and the discount rate $r_{*,t+i}^{j,e} = r_{t+i}^{j,e} + 0.0577$ for $i = 0, 1, 2, ..., I$.

Step 2(d): The present discounted value of each of the individual J streams of dividends is now taken in accordance with Equation 4. Note that for each of the j streams the present value price so calculated for that stream is the ex post rational price for the j^{th} economy; the price a rational investor would pay for the stock if she knew for certain that the j^{th} economy would obtain. We can therefore call the j^{th} present value $P_t^{X,j,e}$, where the superscript X denotes ex post rational.

[13] 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.

In considering these prices, note that according to Equation 4 the stream of discount and dividend growth rates should be infinitely long, while in our simulations we extend the stream only a finite number of periods, I. Since the ratio of gross dividend growth rates to gross discount rates – i.e., the ys in Equation 6 – are less than one in steady state, the individual product elements in the infinite sum in Equation 6 (equivalently Equation 4) eventually converge to zero as I increases. Indeed, this convergence to zero is exactly what is required for the absence of price bubbles! We therefore set I large enough in our simulations so that the truncation does not materially effect our results. As mentioned previously, we found that setting $I = 400$ (years) accomplishes this conservatively. That is, the discounted value today of a dividend payment received 400 years in the future is essentially zero. Also note that the steps above are required to produce P_t^e, D_t^e, g_{t-1}^e, and $r_{*,t}^e$ for $e = 1, ..., E$; the intermediate terms superscripted with a j are required only to perform the numerical integration that yields P_t^e.

Step 2(e): Perform Steps 1(a)-(c) and 2(a)-(d) for $t = 1, ...T$, rolling out E independent economies for T periods. The length of the time series T is chosen to be 47 to imitate the 47 years of annual data we have available from the S&P 500 from 1952-1998.[14] This produces D_t^e, r_t^e, and P_t^e, for $t = 1, 2, 3, ..., T$ and $e = 1, 2, 3, ..., E$. We can also construct the market returns for these economies,

$$R_t^e = (P_{t+1}^e + D_{t+1}^e - P_t^e)/P_t^e$$

and the equity premium, π^e, that agents in the e^{th} economy would observe. The equity premium satisfies the pricing-kernel condition:[15]

[14] To avoid initial conditions contaminating the simulations, the dividend growth rates and interest rates are simulated for over a hundred periods before dividends and prices are calculated.

[15] As with determining the equity premium for the S&P 500 data, this requires a nonlinear estimation problem to be solved and was accomplished with a grid search. The equity premium is restricted to be positive and thus is set to 0 if $\pi^e < 0$.

$$E\left\{\frac{1+R_t^e}{1+r_t^e+\pi^e}\right\}=1.$$

Exhibit 3 summarizes this procedure by which we calculate market prices for the E economies in Step 2(e).

———— Exhibit 3 goes here ————-

II.B.2. Calculating Fundamental Prices

Now that we have calculated market prices for each of the E economies, we next move to calculate estimated fundamental prices for the same economies based on various fundamentals-estimation procedures, such as the Gordon and DK models.

Step 3: It is useful for future reference to first calculate ex post rational prices for these E economies; these are the prices that obtain by substituting realized interest rates and dividends for their expected values in Equation 2. To do this, Define $P_t^{X,e}$ as the ex post rational price for economy e at time t and treat P_T^e as the truncation point price; i.e., the last price observation in the sample.[16] We will use the realized dividends D_t^e, interest rates r_t^e and equity premium π^e to perform the calculation for each economy $e=1,2,3,...,E$. This calculation of ex post rational prices for each of the E economies and $t=1,...,T$ time periods produces E *time series* of ex post rational prices produced by our simulations.

Step 4: We must also form Gordon Model prices for each of these E economies. That is, for each of the $e=1,2,3,...,E$ simulated economies for which market prices were calculated in the previous section, we now calculate a fundamental price using the Gordon model.

Step 4(a): Set $\bar{g}^e=\sum_{t=1}^{T} g_t^e/T$ and $\bar{r}_*^e=\sum_{t=1}^{T} r_t^e/T+\pi^e$, and form the Gordon price

[16] In the literature (e.g., Shiller(1981)), when calculating ex post rational prices, a truncation point is chosen to be either fixed, at say the last observed market price, or to be a moving point 20 or more years ahead of the date for which we are calculating the price. With a fixed date, the ex post price converges to the market price at the terminal date.

$$P_t^{G,e} = D_t^e \left[\frac{1+\bar{g}^e}{\bar{r}_*^e - \bar{g}^e} \right]$$

.

Step 4(b): In addition to the classic Gordon model, there have been a variety of extensions to make the Gordon model more realistic by allowing dividend growth rates to vary over time. We implement the trinomial dividend models of Yao (1997): the additive Markov model and the geometric Markov model specified in Equation 8 and Equation 9 above.

The additive model has us estimating for each $e = 1, 2, 3, ..., E$ economy the average absolute value of the change in the dividend payment, $\Delta^e = \sum_{t=2}^{T} |D_t^e - D_{t-1}^e|/(T-1)$, the proportion of the time the dividend increases, $q^{e,u}$, and the proportion of the time the dividend decreases, $q^{e,d}$, in order to form the price estimate

$$P_t^{ADD,e} = D_{t-1}^e/\bar{r}_*^e + \left[1/\bar{r}_*^e + (1/\bar{r}_*^e)^2\right] \left(q^{e,u} - q^{e,d}\right) \Delta^e.$$

The geometric model has us estimating for each economy $e = 1, 2, 3, ..., E$ the average absolute value of the percentage change in the dividend payment, $\Delta^{e,\%} = \sum_{t=2}^{T} |(D_t^e - D_{t-1}^e)/D_{t-1}^e|/(T-1)$, to form the price estimate

$$P_t^{GEO,e} = D_{t-1}^e \left[\frac{1 + (q^{e,u} - q^{e,d})\Delta^{e,\%}}{\bar{r}_*^e - (q^{e,u} - q^{e,d})\Delta^{e,\%}} \right]$$

Step 5: Finally, we also consider the Donaldson and Kamstra (1996) procedure, which focuses on the ratio of dividend growth to discount rates, $log(y_t) \equiv log((1+g_t)/(1+r_t+\pi))$, rather than on dividend growth and interest rates separately. The DK procedure calls for evolving many possible streams of ys into the future with a Monte Carlo simulation and then taking the present value as in Equation 10. Donaldson and Kamstra (1996) argue that y_t is better behaved than g_t or discount rates alone and that it makes more sense to forecast y_t since this is the object of investor interest in Equation 6. In other words, investors care about the ratio of gross dividend growth to the gross discount rate, not about each variable individually, so it makes more sense to forecast the ratio y. The DK procedure is described below and summarized in Exhibit 4.

———— Exhibit 4 goes here ————

Step 5(a): To implement the DK procedure, we start by estimating a model for $log(y)$ that captures its dynamic evolution over time. For each of the $e = 1, 2, 3, ..., E$ simulated economies for which market prices were calculated in the previous section, we now use the Bayesian Information Criterion to select the "best" model from the set of ARMA(p,q) models, (p,q)=(1,0), (1,1) and (2,0) for $log(y_t)$. We find that (p,q)=(1,0) is usually, but not always, chosen. For each of the E economies we then use the estimated model to make conditional mean forecasts $\widehat{log(y_t^e)}$, $t = 1, ..., T$, conditional on only data observed before period t for that economy. We also estimate the sample standard deviation ($\hat{\sigma}^e$) of the residual of the model within each economy. With $E = 1000$ economies we would do this step 1,000 times, once for each of the $e = 1, ..., E$ economies.

Step 5(b): Now simulate discounted dividend growth rates for each of the $e = 1, 2, 3, ..., E$ simulated economies. That is, produce $log(y^e)$ that might be observed in period t in economy e given what is known at period $t - 1$ in economy e. To do this for a given period t and economy e, simulate a population of J independent possible shocks (draws from a normal distribution,[17] mean 0 and standard deviation equal to $\hat{\sigma}^e$), which we label $\epsilon_t^{j,e}$, $j = 1, ..., J$, and add these shocks separately to the conditional mean forecast $\widehat{log(y_t^e)}$ from Step 5(a) so as to produce $log(y_t^{j,e}) = \widehat{log(y_t^e)} + \epsilon_t^{j,e}$, $j = 1, ..., J$. This is a simulated cross-section of J possible realizations of $log(y_t^e)$ considered at time $t - 1$ for economy e, i.e. different paths that economy e may take next period. Note here that we are performing a

[17] The regression error from time series models of $log(y)$, formed with the S&P 500 and T-bill data, are normally distributed, leading us to this choice of distribution for the shocks here. Limited experiments with the Monte Carlo $log(y)$ series indicates the regression error invariably appears to be normally distributed. As y is a ratio of log normal random variables, this is not surprising.

Monte Carlo simulation on each of the e economies that were themselves generated by a Monte Carlo simulation, so that if we generate $E = 1000$ economies in Steps 1-2 above, and $J = 1000$ economies in Step 5, we are in total generating $1,000,000$ economies for the DK simulation (which is precisely what we have done for the investigations in this chapter).

Step 5(c): Next, for each economy e, use the estimated model from Step 5(a) to make the conditional mean forecast $\widehat{log(y_{t+1}^{j,e})}$, conditional on only the j^{th} realization for period t, $log(y_t^{j,e})$ and $\epsilon_t^{j,e}$, and the data known at period $t-1$, and simulate a population of J independent shocks $\epsilon_{t+1}^{j,e}$, $j = 1, ..., J$ as in Step 5(b) to form $log(y_{t+1}^{j,e})$.

Step 5(d): Repeat Step 5(c) to form $log(y_{t+2}^{j,e}), log(y_{t+3}^{j,e}), ..., log(y_{t+I}^{j,e})$ for each of the J economies, where I is the number of periods into the future the simulation is run ($I = 400$ in our simulations, as explained above). Then form the simulated ex post rational price, $P_t^{s,j,e}$, where

$$P_t^{s,j,e} = D_t^e \left(y_t^{j,e} + y_t^{j,e} y_{t+1}^{j,e} + y_t^{j,e} y_{t+1}^{j,e} y_{t+2}^{j,e} + \cdots \right); \; j = 1, ..., J,$$

corresponding to the J possible economies based on the previously simulated economy e.

Step 5(e): Calculate the DK fundamental price for each time period $t = 1, ..., T$ and each economy $e = 1, 2, 3, ..., E$:

$$P_t^{DK,e} = \sum_{j=1}^{J} P_t^{s,j,e} / J.$$

The DK procedure outlined above is represented diagrammatically in Exhibit 4 above. For the experiments in this chapter we set $E = 1000$, $J = 1000$, and $I = 400$.[18]

[18] Less than 2% of the simulations (economies) yielded DK models that were not stable, or that produced rollouts that were not stable, and were excluded from the analysis below. These economies were not remarkable otherwise, and the majority of the experimental results we present are qualitatively unchanged whether these simulations are included or not.

II.B.3. Sensitivity of the Monte Carlo Results

Careful analysis of any Monte Carlo simulation must include a discussion of the simulation error itself. In a world of unlimited resources, the simulation error can be driven down to negligible scales by increasing the number of replications, in our case increasing the number of simulated economies, E, from 1,000 to several million, and increasing the fans, J, used in the calculation of each economy's market price from 1,000 to several million. This chapter's Monte Carlo experiment involves a simulation of a numerical integration (to calculate market prices) as well as the method of Donaldson and Kamstra (1996) which is also a numerical integration (to calculate the DK prices). The scale of the simulation quickly reaches frightening proportions as we increase the number of replications (economies), E, or fans, J, in the numerical integrations. The choice of 1,000 economies, each stretching out for 47 years, and then the choice of 1,000 fans used to calculate the market prices for each year of each economy, led to roughly one month of CPU time on an SUN UltraSparc II 400 Mhz machine, the practical limit for this experiment given competing demands for this resource.

To determine the simulation error, we must conduct 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 1,000 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 thousand price estimates would be identical. With the number of fans, J, equal to 1,000 we find that the standard deviation of the simulation error is only 0.28 % of the price, which is sufficiently small as to not be a source of concern for our study.[19]

[19]We further investigated the sensitivity of selected results by increasing the number of simulated

III. The Results of the Monte Carlo Experiments

We now explore results based on our simulated market economies and the fundamental prices calculated from them. The first set of results focus on summary statistics for returns, dividend-price (D/P) ratios, dividend growth rates, interest rates and risk premia produced by our simulated economies as compared to data from real-life financial markets. The second set of results focuses on tests of bubbles in the simulated economies, as well as on some related statistics.

Recall that no experiment presented here is calibrated to actual S&P 500 *prices*. All experiments are calibrated only to dividend growth rates and discount rates, and then prices are calculated based on the underlying fundamental factors. Great care was taken to ensure that our experiments faithfully replicated the mean, variance, cross-correlation and autoregressive structure of log dividend growth rates and log discount rates. If S&P 500 prices contain bubbles, we would expect that S&P 500 returns and price/dividend ratios would be excessively volatile relative to our simulated economies, which do not themselves contain bubbles. Furthermore, since we have no bubbles in our simulated economies, the properties of the fundamental price estimates and the ex post rational price under the null of no bubbles can be obtained from our simulations. This allows us to investigate the size properties of tests for bubbles based on these bubble-free price estimates.

——— Table 1 Goes here ———

Recall from the discussion above that we have 47 years of data from the S&P 500 and that each of our 1000 simulated economies therefore also produce 47 years of data. For each of the 1001 economies under investigation (1 actual economy and 1000 simulated economies) we calculate the mean, standard deviation, skewness and kurtosis of a variety of interesting

economies, E, or the number of fans, J, to as large as 10,000 (restricting to two years the number of periods each economy was followed out), but found no qualitative changes to our experimental results. Of course, this did reduce the simulation error by a factor of square root 10, but was not of much practical gain since the simulation error was already extremely small.

variables, including dividend growth rates, interest rates, stock returns and dividend-price ratios. For example, we calculate the mean return of the 47 years of the actual S&P 500 which appears in the top left cell of the results shown in Table 1. We also calculate the mean return for each of our 1000 simulated economies, which produces 1000 mean return estimates, the distribution of which is summarized in the remaining cells of the first row of results shown in Table 1 (the X^{th} percentile reports the return for the economy with the X^{th} largest return).

We begin our analysis of the results in Table 1 by considering results for dividend growth rates and discount rates. Given that we calibrated the discount rates and dividend growth rates to the true data, we would expect that discount and dividend growth rates from our simulated economies would appear quite similar to the true market data. And indeed, we see means, standard deviations, skewness and kurtosis for our simulated interest rates and dividend growth rates very nearly identical to the true interest rate and dividend growth rate data.[20] Furthermore, each and every statistic for dividend growth rates and interest rates lies near the center of the distribution of the simulated series. Given that we generated dividend growth rates and interest rates to match the true data this is not surprising and is only presented as a verification check. The true challenge, which may provide some insight into the existence of bubbles in the S&P 500 data, is to see if S&P 500 returns and prices also look like returns and prices from economies that share dividend and interest rate characteristics with the true economy, but have been simulated to have no bubbles.

From the first row of Table 1 we see that, on average, our simulated economies produce mean returns very close to returns from the actual S&P 500. Indeed, the S&P 500's return lies near the middle of the distribution of returns from our simulations. The skew and kurtosis of S&P 500 returns also fall well within the 90% confidence interval formed by the simulated economies. Only the standard deviation of market returns is (barely) outside the

[20]For instance, the S&P 500 dividend growth rates have a slight positive skew of 0.257, while the 5 and 95 percentiles of the simulated series of dividend growth rates are -.45 and 0.66 respectively.

90% confidence interval formed by our simulated economies. Price levels themselves can of course not be compared directly since price is a random walk and thus the unconditional mean, standard deviation, skewness and kurtosis do not even exist. Financial studies therefore typically focus on the dividend-price ratio, since prices and dividends are cointegrated. From Table 1 we see is that the D/P ratio of the S&P 500 is similar to our simulated data, with all statistics measured for the S&P 500 data within 90% confidence interval formed by the simulated economies' statistics (with the sole exception of the mean D/P ratio which is outside the 90% confidence interval but within the 98% confidence interval. the 5 percentile range). Again, similar to returns, the S&P 500 D/P ratios do not stand out as wildly unusual, given the behavior of returns we see in no-bubble economies. In sum, the results for dividend growth rates, interest rates, returns and D/P ratios shown in Table 1 suggest that the behavior of the true S&P 500 stock index over the past 50 years is not out of line with how we might expect a financial market to behave if that market priced stocks in a rational, efficient and bubble-free manner, given interest rate and dividend fundamentals. What other studies have argued to be anomalies in price behavior, such as volatile return series, seem in our study to be entirely consistent with prices based on dividend and interest rate fundamentals. The key difference between our investigation and previous work is our calibration to the dividend and interest rate characteristics of the economy with no simplifying assumptions when simulating market prices, returns and various statistics and their distributions derived from these prices and returns. Conversely, previous studies, even those that have used simulation techniques, can only make (educated) guesses at what the empirical distribution of various tests and statistics might be without the simplifying assumptions imposed. To the extent that these guesses and simplifying assumptions (e.g., independent dividend innovations) are inaccurate, or to the extent that small sample properties of the financial time series are not accounted for, misleading results can be obtained and "anomalies" can be found where none truly exist. Thus the approach we adopt leads to a conclusion opposite to many studies in the literature, starting with Shiller (1981), which conclude that stock prices are excessively volatile relative to fundamentals.

———— Figure 1 goes here ————

To further investigate the summary statistics reported in Table 1, we provide a series of figures which show the bivariate distributions of the economy-by-economy statistics that went comprise Table 1. The S&P 500 economy could conceivably be relatively usual compared to the simulated economies, statistic by statistic, but have a very unusual combination (pairing) of statistics, which would become obvious from a bivariate plot of these statistics for all the economies (e.g., the standard deviation of returns and the mean of returns for the S&P 500 might each appear similar to simulated data, but the mean-standard deviation combination seen in the S&P 500 data might be very different than the mean-standard deviation combinations produced by the simulations). The plots in Figures 1 through 5 display data points for each of our simulated economies and for the S&P 500, with the actual S&P 500 data represented by a large black dot and the simulated data by smaller circled points (which can, on occasion, appear in the plots to be small black dots because of several circles partially overlaying each other).[21]

Figure 1 presents standard deviations versus means in a group of four panels for returns (Panel A), D/P ratios (Panel B), dividend growth rates (Panel C) and interest rates (Panel D). We would expect the dividend growth rate and interest rate plots to have the actual S&P 500 data in the center of the cloud of points associated with the simulated economies since we calibrated the experiments to these two quantities, and this is indeed what we observe in Figure 1, Panels C and D. As for returns in Panel A, we see that economies generated by calibrating to only discount rates and dividend growth rates produce mean/variance pairs for returns that do not look terribly dissimilar to the S&P 500 return data. The plot of the D/P ratios means versus standard deviations in Panel B also shows

[21]For presentation, data points which completely or very close to completely overlay each other were condensed down to a single point.

that the first two moments of the S&P 500 data are not entirely out of the bounds of possibility, although the S&P 500 data do stray to the very edge of the joint distribution of means and standard deviations for the D/P ratios from the simulated economies. In other words, prices in the S&P500 data seem higher relative to dividends than we typically see in our simulated economies. This suggests that, while the S&P500's D/P ratio is not low enough to claim with any degree of confidence that current market prices contain a bubble of irrational exuberance, our evidence does suggest that actual market prices do seem rather high relative to what we observe in our simulated economies. In other words, S&P500 stock prices may indeed be suspiciously high as some people suggest, but they are not so high as to allow us to conclude that they are not being driven by fundamentals.

———— Figure 2 goes here ————

Turning our attention back to Table 1, in the last row of Table 1 we present equity premia from the true data and equity premia from our simulated economies. Note that since we have only one equity premium value per economy we cannot present a standard deviation, skew or kurtosis measure for the S&P 500 data; there is only the one equity premium estimate of 5.77%. Similarly, in each of the simulated economies we have a single measure of the equity premium and thus cannot present a standard deviation, skew or kurtosis measure for the simulated equity premia either. However, we do see from Table 1 that the risk premia possible in a simulated economy, with dividend and interest rate processes set to mimic actual dividend and interest rate processes, can range from 0 to over 8%, even if the true equity premium is below 6%. This suggests that the equity premium we see in the actual S&P 500 data could be an imprecise estimate of the market's true underlying risk premium and leads us to wonder whether findings of so-called anomalies, such as the Mehra and Prescott (1985) "equity premium puzzle", are really puzzling, or whether such findings are simply small sample flukes.[22]

[22] For an initial identification of the equity premium puzzle see Mehra and Prescott (1985). This seminal work spawned a significant body of research on the equity premium puzzle, much of which is reviewed in

To further examine the equity premium issue, we plot in Figure 2 the relationship between stock returns and the equity premium. Recall that the actual S&P 500 data are represented by a large black dot and the simulated data by smaller circled points. Although the high equity premium of the last half-century has been much discussed, and argued to be excessive, we see that the distribution of possible equity premia and mean market returns over a 47 year period is quite spread out and that the S&P 500's experience is not terribly unusual, at least if you allow discount rates and dividend growth rates to be autocorrelated random variables as we do. In fact, in a bubble-free economy with fundamentals (i.e. dividend growth rates and discount rates) calibrated to match the true economy, it is possible to observe measured equity premia very nearly equal to mean returns, even when we would expect that over hundreds of years the mean return would converge to 11% (the mean return of the S&P 500 over the past 47 years) and the mean equity premium to 6% (the equity premium of the S&P 500 over the past 47 years). The variance of the S&P 500 discount rates and market returns is so large that there is very little we can draw from the so-called equity premium puzzle in the short time periods available to us in the true economy.

———— Figure 3 goes here ————

Some authors have argued that the equity premium is "too large" not because stock returns are too high, as originally argued, but rather because riskfree interest rates are too low.[23] Figure 3 therefore presents mean interest rates versus the equity premia. Again we find the experience of the true data over the last 47 years to be unremarkable relative to our simulated economies. It is not uncommon to find in our simulations very low mean interest rates and high risk premia. In fact, low average interest rates are virtually a prerequisite to finding high equity premia, as the sum of the two is restricted to add up (approximately) to the mean return, and the mean return rarely strays above 14% over a

Campbell, Lo and MacKinlay (1997).

[23]For an excellent review, see Kocherlakota (1996).

47 year period, according to our simulations.

——— Figure 4 goes here ———

For completeness, in Figure 4 we present mean returns versus interest rates. This combination shows the actual S&P 500 data to have a relatively high return relative to the interest rate, but even in this case the S&P 500 data is within the envelope of the simulated economies.

For all the data and statistics we explore, including results on returns, D/P ratios, equity premia, etc., the S&P 500 data does not seem to be wildly unusual among economies we have generated. These economies have been generated to be bubble-free and calibrated to the S&P 500 dividend and discount rate experience. Taken together this evidence therefore suggests that the S&P 500 index does not contain price bubbles. Given this evidence that the stock market does not contain bubbles, why is the financial econometric literature replete with tests for bubbles that in general support the argument that stock prices are excessively volatile? It is to this question that we now turn.

——— Table 2 Goes here ———

Table 2 presents the results of four common statistical tests for stock price bubbles, presented in rows one through four of the table, with the columns indicating the pricing method (the the ex post rational price and four different fundamental price estimates) the test was based on. Each test for bubbles was applied at the 10% significance level for each pricing method (where applicable). In Table 2 we present the actual rejection rates across our simulated economies for each of these tests and for each of the fundamental price estimation techniques. If each test is well-behaved – i.e., has proper size – and if these fundamentals estimation techniques are reliable guides to the true fundamental value, then we would expect approximately 10% rejection rates in each entry for the first four rows of the table, as all these tests are being applied under the null of no bubbles.[24]

[24]The rejection rate we calculate with these simulations is a random variable, centered on 10% with

From the first four rows of Table 2 we see that the most commonly applied tests for bubbles, the Camerer test and the two MRS tests, are extremely unreliable. If the tests have correct size, the p-value entries in Table 2 would all be 0.10. However, in all but one case the empirical p-values far exceed their theoretical values, which implies that these tests typically grossly over-reject. In other words, common tests for stock price bubbles find bubbles where none actually exist. Indeed, the Camerer test finds bubbles 100% of the time using ex post and Gordon fundamental price series and 59% of the time using the DK fundamental price series, even though the market price series does not contain bubbles. Only the Campbell-Shiller test procedure (the Dicky-Fuller unit root test) comes close to having correct size.

The last two rows of Table 2 report results on the fundamental pricing error, defined as the fundamental price minus the simulated market price, all divided by the simulated market price. Of all the fundamentals estimation techniques, the DK procedure is the most reliable forecaster of market prices, with the smallest forecast error variance and among the smallest mean errors. The superiority of the DK procedure over other simpler methods is not very great, however. All the fundamental price estimation methods we employ suffer from the same basic problem, which is the difficulty of estimating key parameters (the mean dividend growth rate and discount rate in the case of the basic Gordon growth method, the parameters of the ARMA process for discounted dividend growth rates in the case of the DK method, etc.). Clearly, based on the results of Table 2, there is scope for improvement in fundamentals estimation techniques.

——— Figure 5 goes here ———

Finally, we explore derived statistics from our experiments, including statistics such as the sample autocorrelation at one lag in market stock returns and the ARCH effect in stock returns. The fashion in which our data are constructed should lead to returns that are

a standard deviation derived from the binomial distribution, equal to the square root of 0.10 times 0.90 divided by the number of simulated economies, or 0.0095 approximately. Thus we would expect the rejection rates to differ from 10% by no more that approximately 0.02 at the 5% level of significance.

independent and identically distributed, but the distribution of the sample autocorrelation at one lag and the ARCH(1) coefficient estimate may not be well approximated with the usual asymptotic distributions. Figure 5 presents the universe of pairs of AR(1)-ARCH(1) coefficient estimates from our simulated economies' time series of returns and also plots the coefficient estimates from the S&P 500 returns data. From Figure 5 we see that, once again, the S&P 500 data do not appear wildly at odds with the cloud of data points from our simulated economies. Furthermore, we see from our simulated economies that by random chance we can find economies that exhibit extremely strong annual ARCH effects, with the ARCH(1) coefficient achieving values ranging from -0.4 to 0.6, and equally strong autoregressive effects, with positive and negative magnitudes of roughly 0.50 on the AR(1) coefficient, even though neither of these effects are actually present in the data.

IV. Conclusions

The purpose of this chapter has been to use advanced simulation and Monte Carlo techniques to learn more about the behavior of fundamental stock prices and the properties of statistical tests commonly used to test for price bubbles in stock market data. We produce simulated stock price series based on dividend growth and discount rate fundamentals so as to simulate rational, efficient and bubble-free economies which share dividend and interest rate characteristics with the true economy. While much conventional wisdom suggests that stock prices are "too volatile" and may be influenced by the "irrational exuberance" of traders, our evidence does not support this conclusion. On the contrary, we find that actual stock prices do not behave significantly differently than our simulated prices from rational, efficient and bubble-free economies. We also find that traditional tests for excess volatility and bubbles over-reject the null of no bubbles in samples of the size traditionally employed in the literature. Indeed, most tests we investigate find overwhelming evidence of price bubbles even though we have constructed the data to have no bubbles. In other words, we demonstrate that traditional tests for bubbles frequently find bubbles in data that do not in fact contain bubbles.

The key difference between our investigation and previous work is our calibration to the dividend and interest rate characteristics of the economy with no simplifying assumptions when simulating market prices, returns and various statistics and their distributions derived from these prices and returns. Conversely, previous studies, even those that have used simulation techniques, can only make (educated) guesses at what the empirical distribution of various tests and statistics might be without the simplifying assumptions imposed. To the extent that these guesses and simplifying assumptions (e.g., independent dividend innovations) are inaccurate, or to the extent that small sample properties of the financial time series are not accounted for, misleading results can be obtained and "anomalies" can be found where none truly exist. While more work remains to be done in this area, we believe that the results of this chapter at least raise some concerns with some previously employed approaches and suggest an avenue for further investigation.

References

Bollerslev, T. and R.J. Hodrick, 1995, Financial Market Efficiency Tests, in :M.H. Pesaran and M.R. Wickens, eds., Handbook of Applied Econometrics, Vol. 1 (Blackwell Publishers, Oxford) 415-458.

Brooks, R. and B. Helms, 1990, An N-stage Fractional Period, Quarterly Dividend Discount Model, The Financial Review 25, 651-657.

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

Campbell, J.Y. and R.J. Shiller, 1987, Cointegration Tests of Present Value Models, Journal of Political Economy 95, 1062-1088.

Campbell, J.Y., A.W. Lo and A.C. MacKinlay, 1997, The Econometrics of Financial Markets (Princeton University Press, Princeton, NJ).

Campbell, J.Y. and R.J. Shiller, 1998, Valuation Ratios and the Long-Run Stock Market Outlook, Journal of Portfolio Management 24, 11-26.

Camerer, C., 1989, Bubbles and Fads in Asset Prices, Journal of Economic Surveys 3, 3-38.

Cochrane, J., 1992, Explaining the Variance of Price-Dividend Ratios, Review of Financial Studies 5, 243-280.

Dickey, D.A. and W.A. Fuller, 1981, Likelihood Ratio Statistics for Autoregressive Time Series with a Unit Root, Econometrica 49, 1057-1072.

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

Donaldson, R.G. and M.J. Kamstra, 2000, Forecasting Fundamental Stock Price Distributions, SFU manuscript.

Engle, R.F., 1982, Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation, Econometrica 50, 987-1008.

Fuller, R.J. and C. Hsia, 1984, A Simplified Common Stock Valuation Model, Financial Analysts Journal 40, 51-54.

Gastwirth, J.L. and H. Rubin, 1971, The Behavior of Robust Estimators on Dependent Data, Purdue University Statistics Mimeo Number 197.

Gordon, M., 1962, The Investment, Financing and Valuation of the Corporation (Irwin, Homewood IL).

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

Hull, J., 1993, Options, Futures and Other Derivative Securities (Prentice-Hall, Englewood Cliffs, NJ).

Hurley, W.J. and L.D. Johnson, 1994, A Realistic Dividend Valuation Model, Financial Analysts Journal 50, 50-54.

Hurley, W.J. and L.D. Johnson, 1998, Generalized Markov Dividend Discount Models, Journal of Portfolio Management 24, 27-31.

Kennedy, P., 1992, A Guide to Econometrics (MIT Press, Cambridge MA).

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

Kocherlakota, N.R., 1996, The Equity Premium: Still a Puzzle, Journal of Economic Literature 34, 42-71.

Lehmann, E.L., 1975, Non-Parameterics (Holden-Day Inc., San Francisco).

Malkiel, B.G., 1963, Equity Yields, Growth, and the Structure of Share Prices, American Economic Review 53, 1004-1031.

Mankiw, N.G., D. Romer and M.D. Shapiro, 1985, An Unbiased Reexamination of Stock Market Volatility, Journal of Finance 40, 667 -687.

Mehra, R. and E. Prescott, 1985, The Equity Premium: A Puzzle, Journal of Monetary Economics 15, 145-161.

Orcutt, G.H., and H.S. Winokur, 1969, First Order Autoregression: Inference, Estimation, and Prediction Econometrica 37, 1-14.

SAS Institute Inc., 1999. SAS Procedures Guide, Version 8, Cary, NC: SAS Institute Inc. Testing for Unit Roots in

Said, E.S. and D.A. Dickey, 1984, Testing for Unit Roots in Autoregressive Moving Average Models of Unknown Order, Biometrika 71, 599-607.

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

Shiller, R., 1981, Do Stock Prices Move to Much to be Justified by Subsequent Changes in Dividends? American Economic Review 71, 421-436.

Shiller, R., 1989, Market Volatility (MIT Press, Cambridge MA).

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

Timmermann, A.G., 1995, Tests of Present Value Models with a Time-Varying Discount Factor, Journal-of-Applied-Econometrics 10, 17-31.

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

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

Yao, Y., 1997, A Trinomial Dividend Valuation Model, Journal of Portfolio Management 23, 99-103.

Table 1

Statistics on the S&P 500 Index,

and Simulated Market Index

Data → / Statistic ↓		S&P 500	Mean of Simulated Data	Percentiles of Simulated Data 1%	5%	95%	99%
Returns	Mean	0.11571	0.10065	0.07300	0.07984	0.12517	0.14214
	Std.Dev.	0.13539	0.09137	0.05361	0.06057	0.13068	0.14969
	Skewness	-0.39741	0.13638	-0.66906	-0.48990	0.71935	1.29074
	Kurtosis	2.75816	3.02994	1.87952	2.09439	4.49553	6.92236
D/P Ratios	Mean	0.03629	0.04872	0.03321	0.03654	0.06724	0.08030
	Std.Dev.	0.00946	0.00955	0.00266	0.00364	0.02062	0.02653
	Skewness	0.44283	0.71078	-0.31723	-0.00336	1.64431	2.11081
	Kurtosis	2.68087	3.12383	1.51714	1.81999	5.77567	8.06559
Dividend Growth Rates	Mean	0.05513	0.05517	0.03731	0.04262	0.06746	0.07141
	Std.Dev.	0.03649	0.03680	0.02651	0.02918	0.04435	0.04828
	Skewness	0.25694	0.09863	-0.71792	-0.45213	0.66111	1.03167
	Kurtosis	3.29200	2.73179	1.82880	2.00100	3.82972	4.51803
Interest Rates	Mean	0.05887	0.06008	0.01588	0.02391	0.11970	0.16291
	Std.Dev.	0.02981	0.02984	0.00626	0.00949	0.06851	0.08978
	Skewness	0.95245	0.86858	-0.16511	0.12500	1.82812	2.30454
	Kurtosis	3.69054	3.42080	1.55605	1.85578	6.74406	8.93784
Risk Premia	Mean	0.05774	0.05102	0.00000	0.01141	0.07599	0.08382

Table 2

Statistics and P-Values from Time-Series Tests

for Bubbles, Fads and Excess Volatility

Fundamental Forecasting Model → Statistic or P-Value ↓	Ex Post	Gordon Growth	Gordon Additive	Gordon Geometric	DK
Camerer test*	1.00000	1.00000	1.00000	1.00000	0.59021
MRS1 test*	.	0.28746	0.05505	0.24873	0.21305
MRS2 test*	.	0.31091	0.32314	0.35474	0.27013
Campbell-Shiller test	0.31600	0.10296	0.02446	0.13354	0.21101
Mean % Error†	0.07682	0.08475	0.17549	0.00187	-0.06629
% Error σ	0.17728	0.17842	0.34676	0.17842	0.15557

*: P-values are from a paired-sample t-test, applied as follows. Let $x_{1,t}$ and $x_{2,t}$ be the two series we are interested in testing for difference in variance, let $\bar{x}_1$ and $\bar{x}_2$ be the sample means of $x_{1,t}$ and $x_{2,t}$ respectively, and define $z_t \equiv (x_{1,t} - \bar{x}_1)^2 - (x_{2,t} - \bar{x}_2)^2$, i.e., the difference in two paired samples. Then, under standard regularity conditions, a test of the null hypothesis that the first moment of z_t is greater than 0 can be interpreted as a test of $x_{1,t}$ having greater variance than $x_{2,t}$ and implemented as a standard paired-sample t-test for the mean of z_t being greater than 0. The test statistic has an asymptotic normal distribution under fairly weak conditions, permitting non-normality as well as some dependence and heterogeneity (see Lehmann (1975) and Gastwirth and Rubin (1971)).

†: Percentage error calculated as price forecast minus market price all divided by price forecast.

Exhibit 1

Diagram of DK Monte Carlo Integration

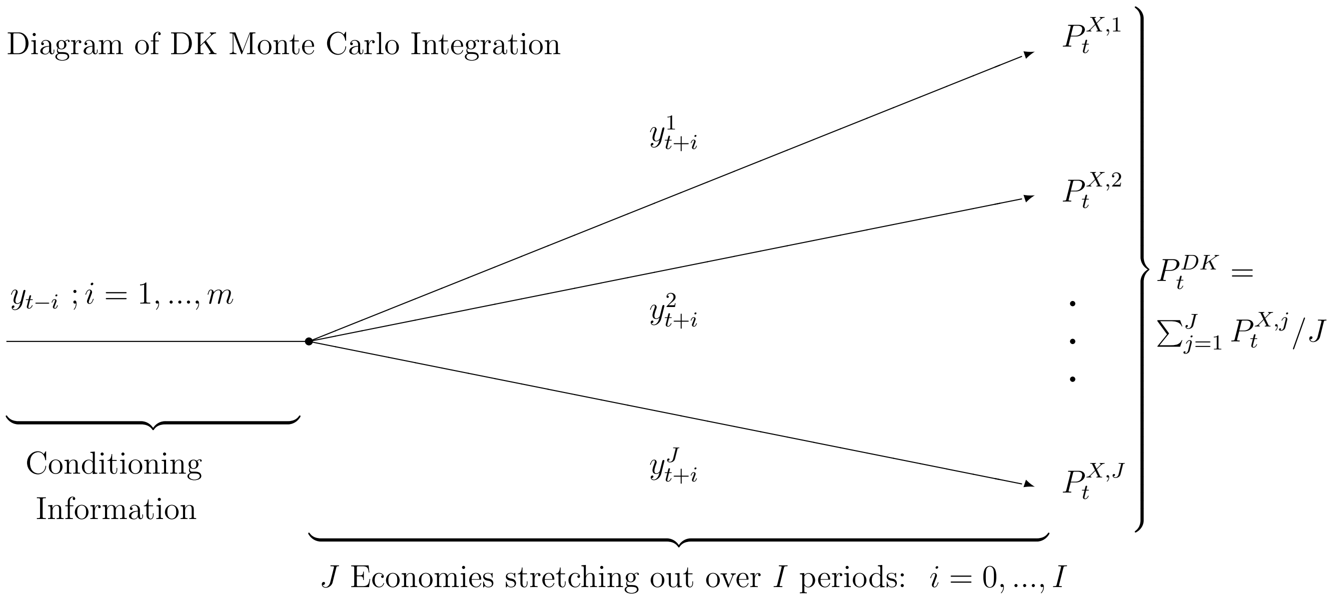

Exhibit 2

Diagram of a Simple Market Price Calculation for the t^{th} Observation of the e^{th} Economy (Steps 1 and 2)

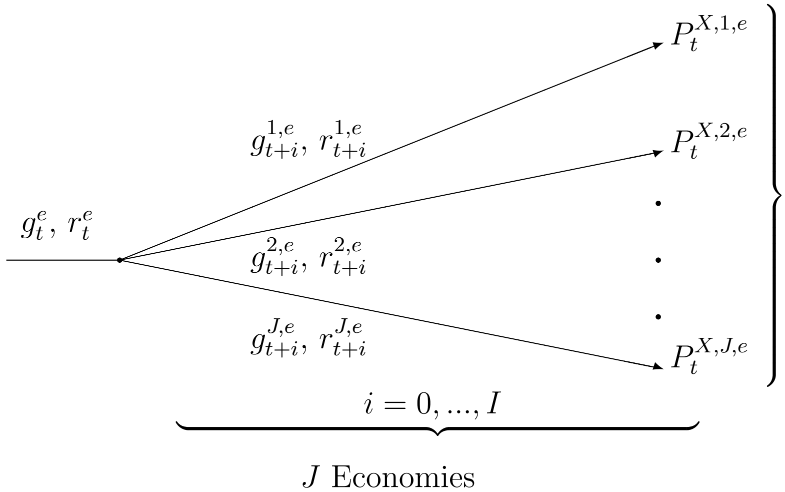

$$P_t^e = \sum_{j=1}^{J} P^{X,j,e}/J$$

Exhibit 3

Diagram of the e^{th} Economy of E Economies and Calculation of the e^{th} Economy Market Price at Time $t+k$ (Step 2(e))

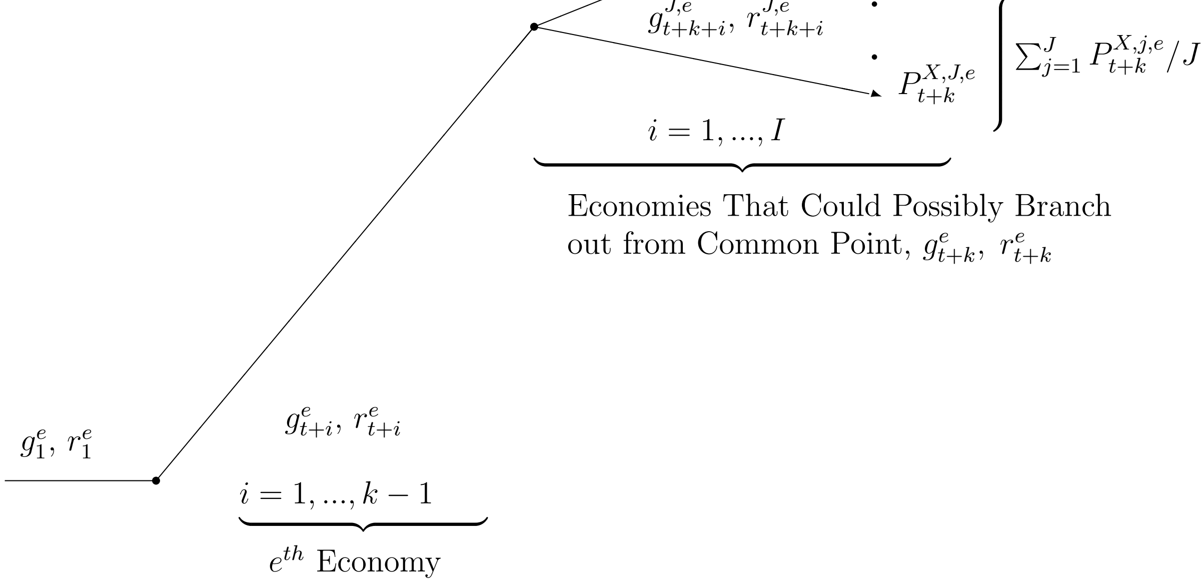

Exhibit 4

Diagram of DK Monte Carlo Integration for the e^{th} Economy

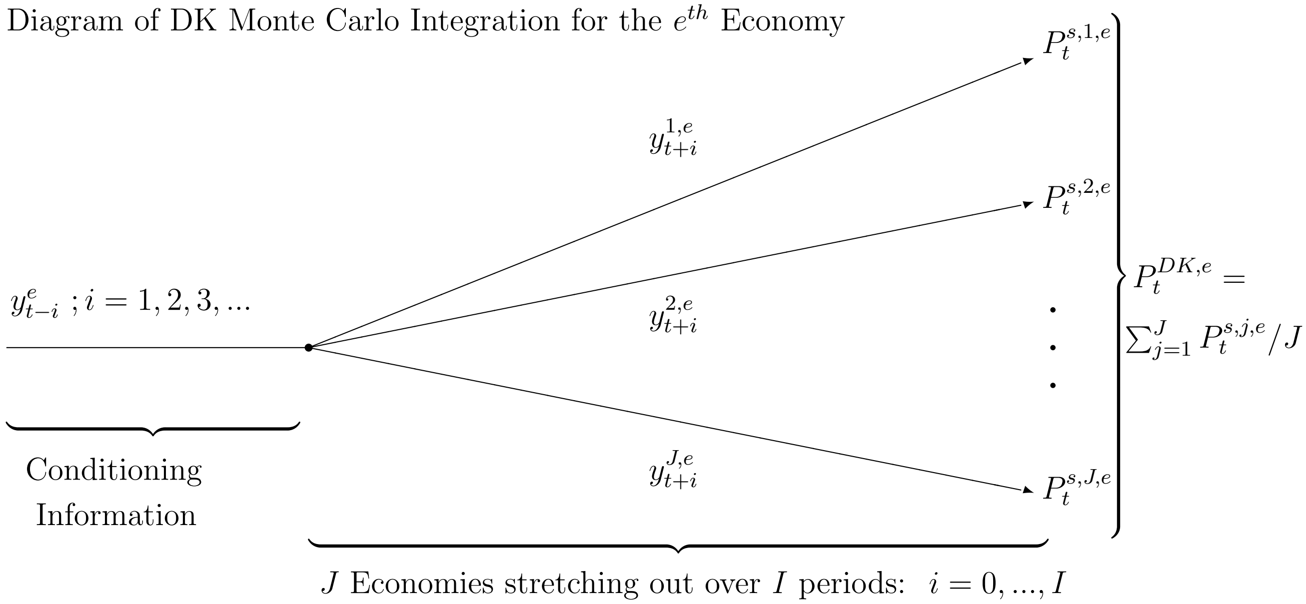

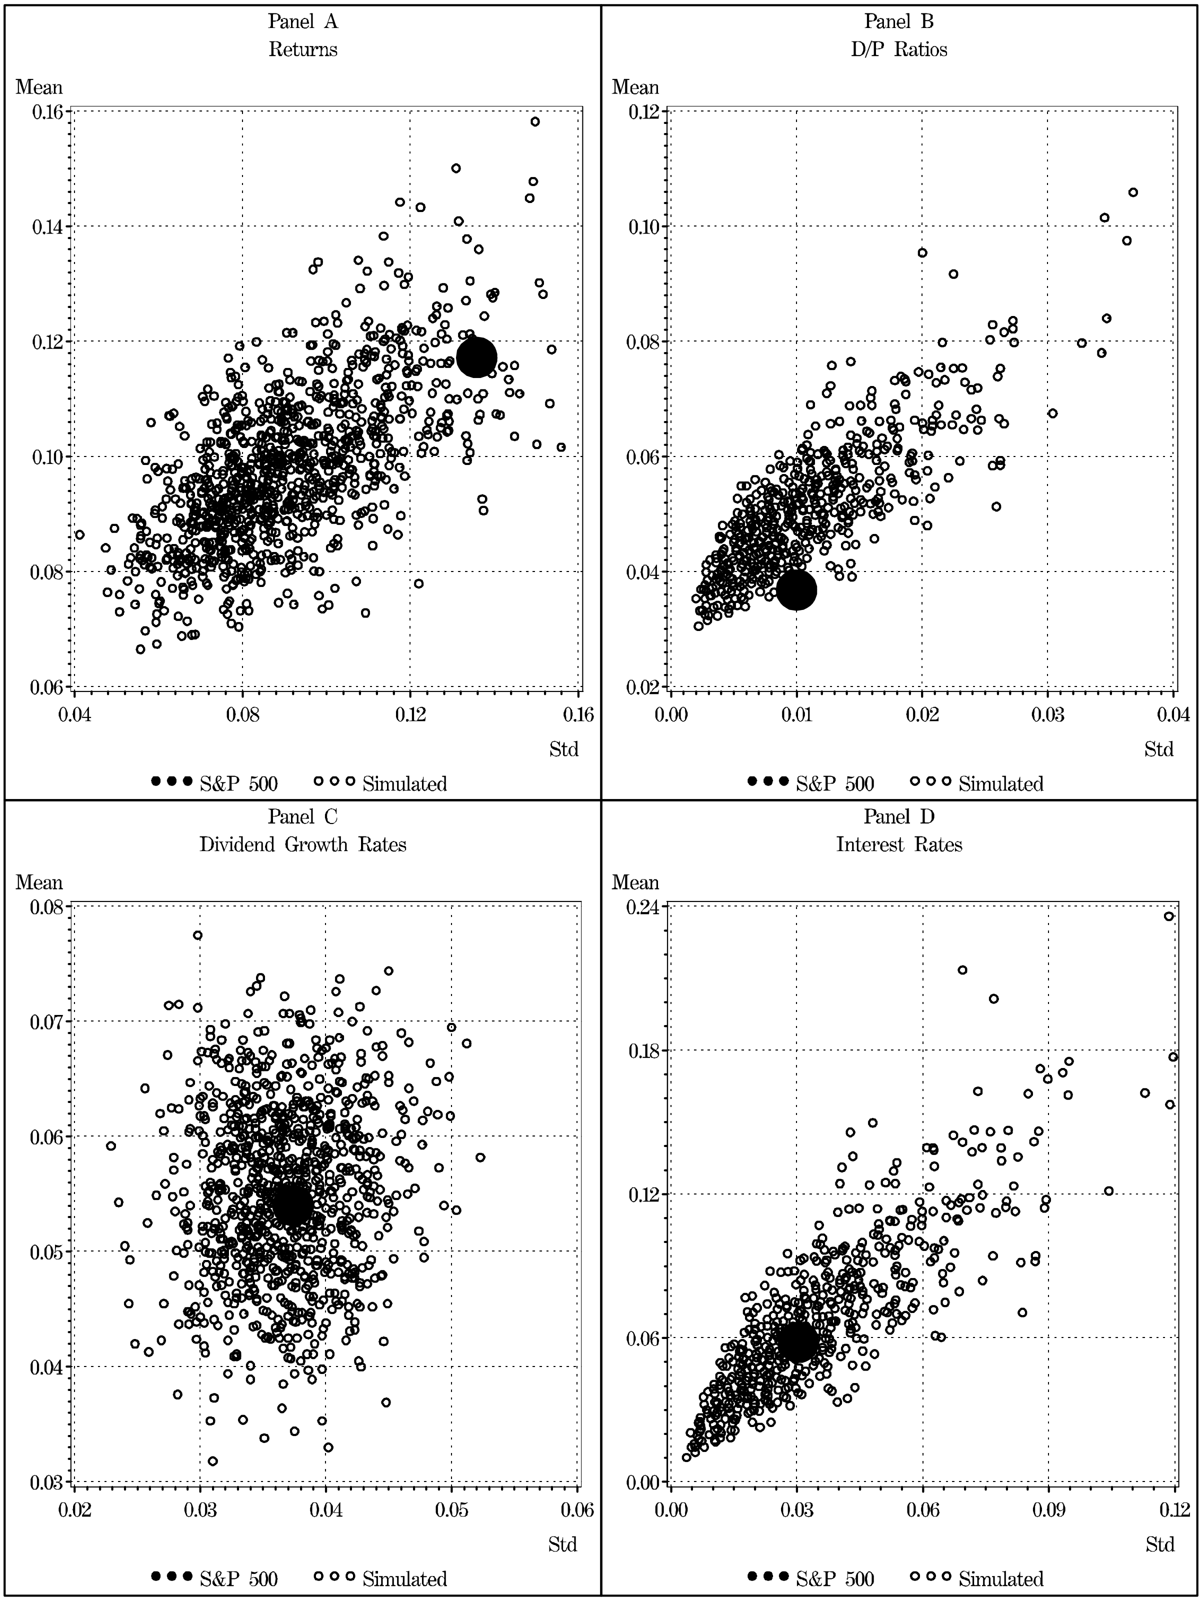

Figure 1: Mean versus Standard Deviation

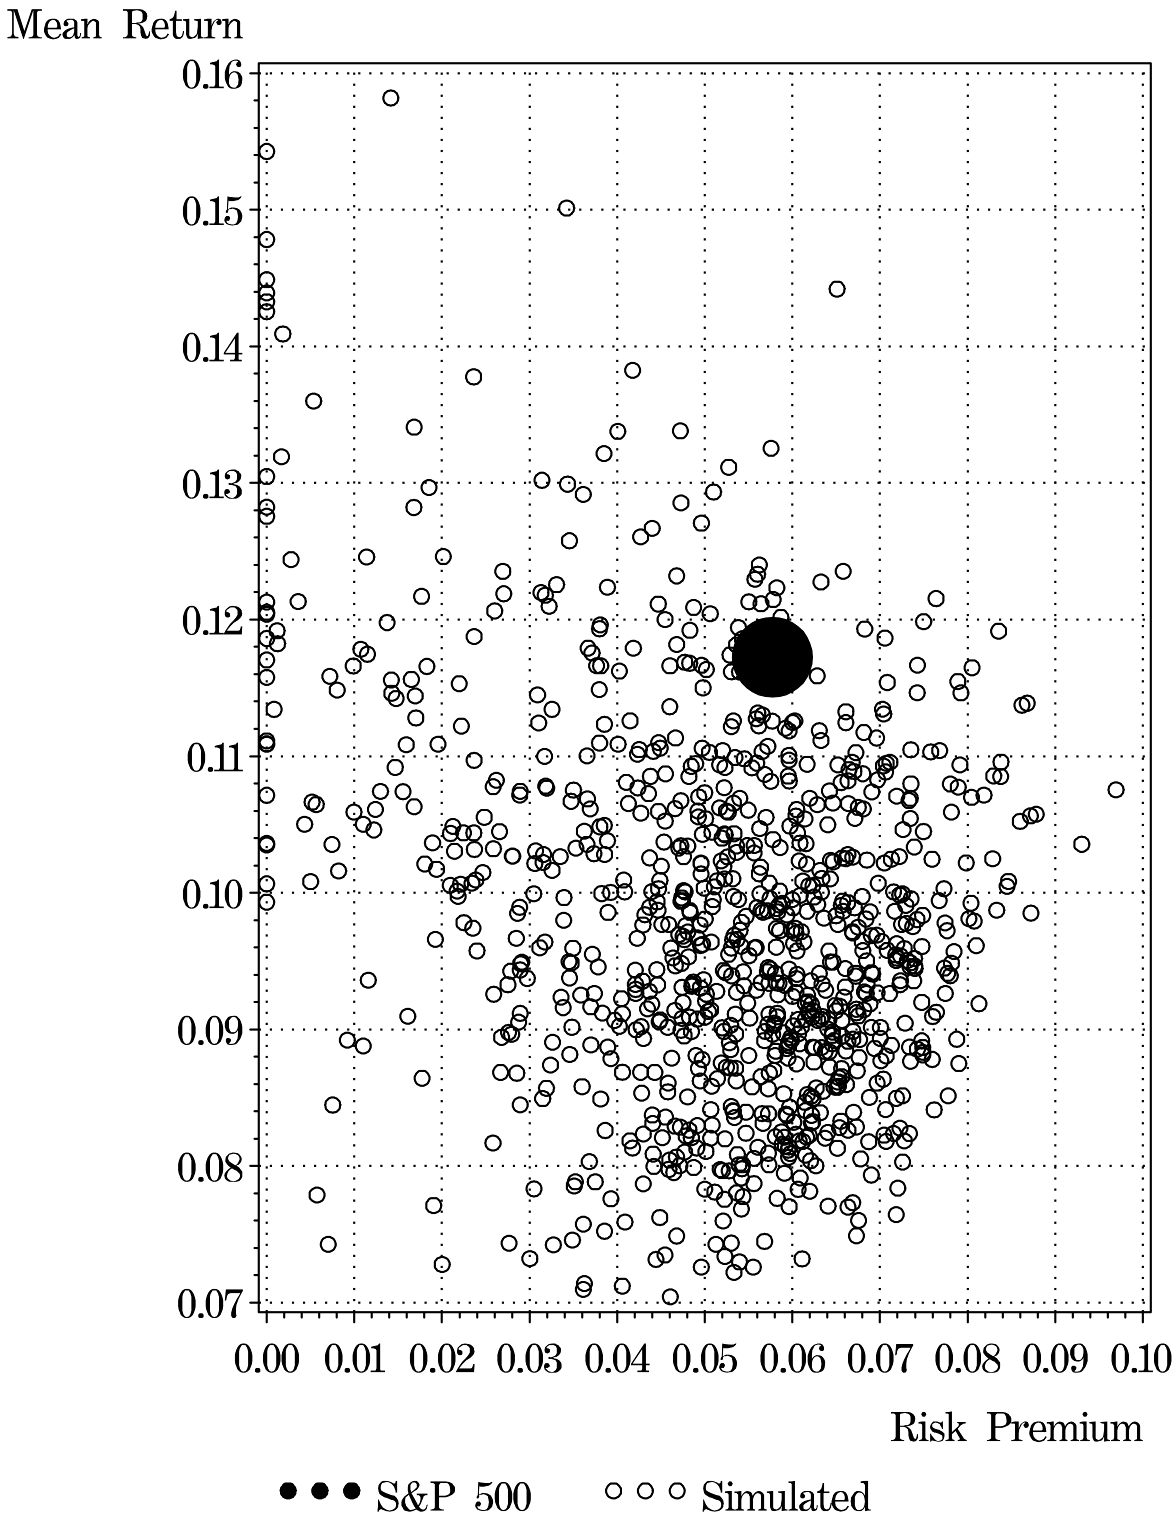

Figure 2: Return Mean versus Risk Premium

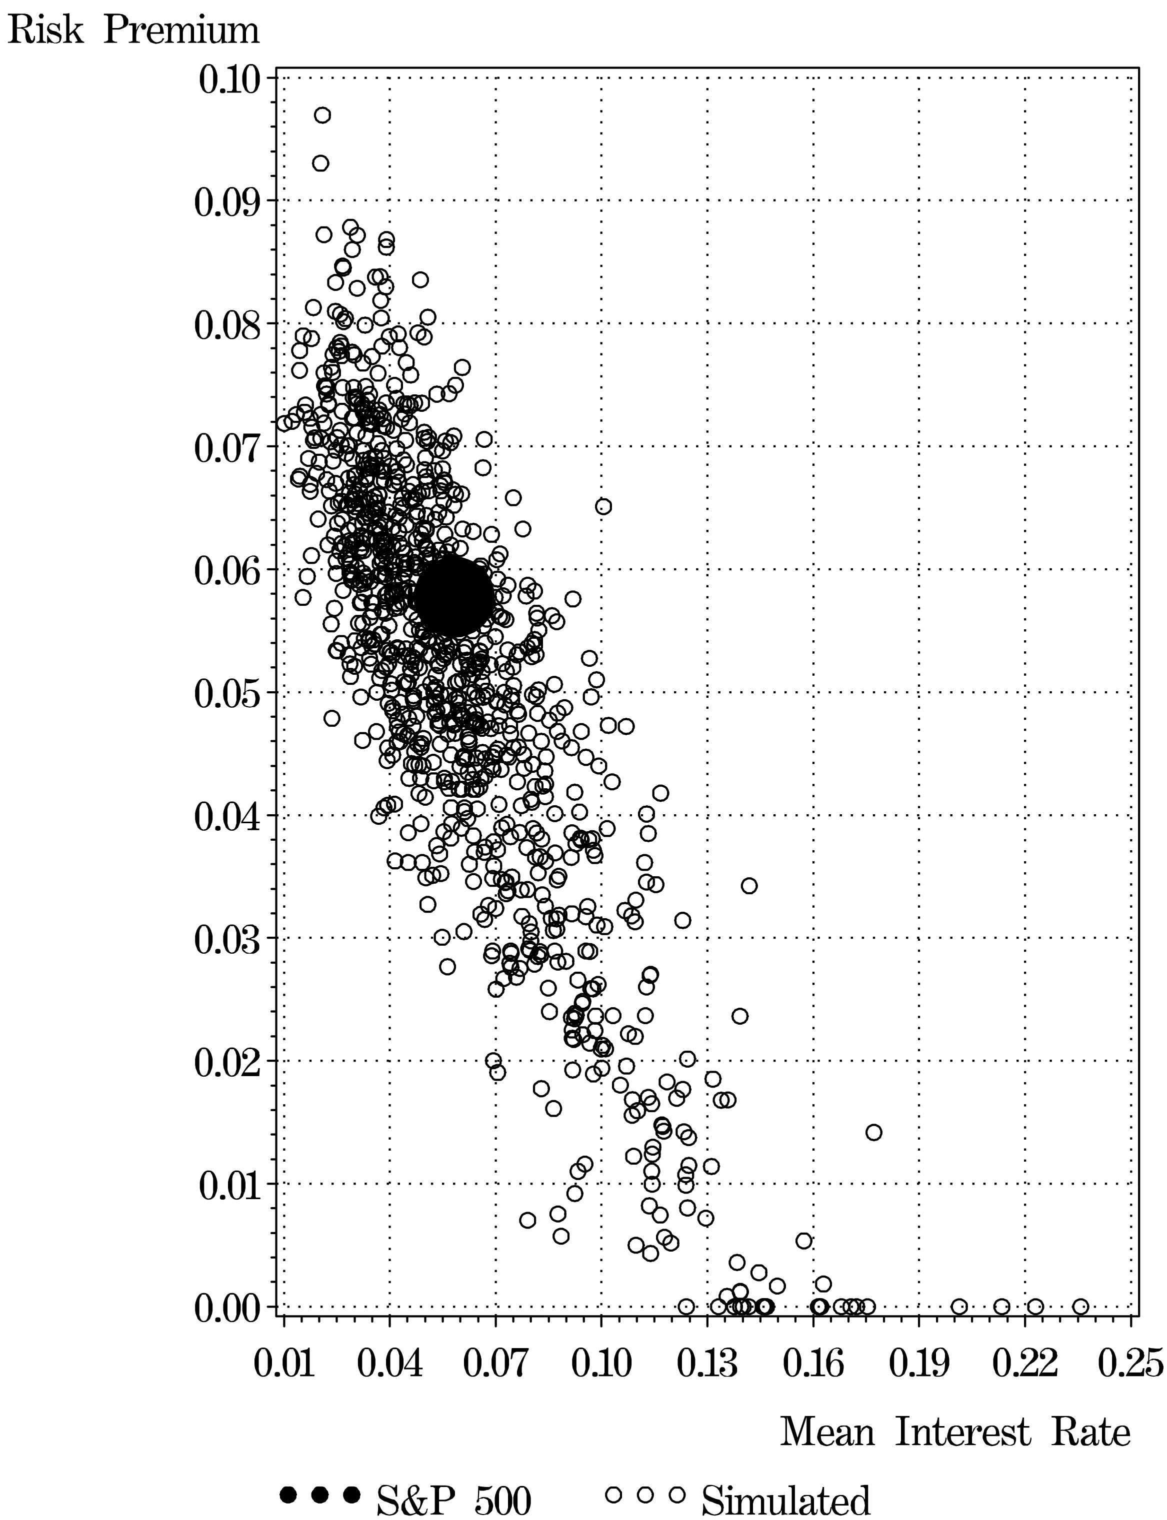

Figure 3: Interest Rate Mean versus Risk Premium

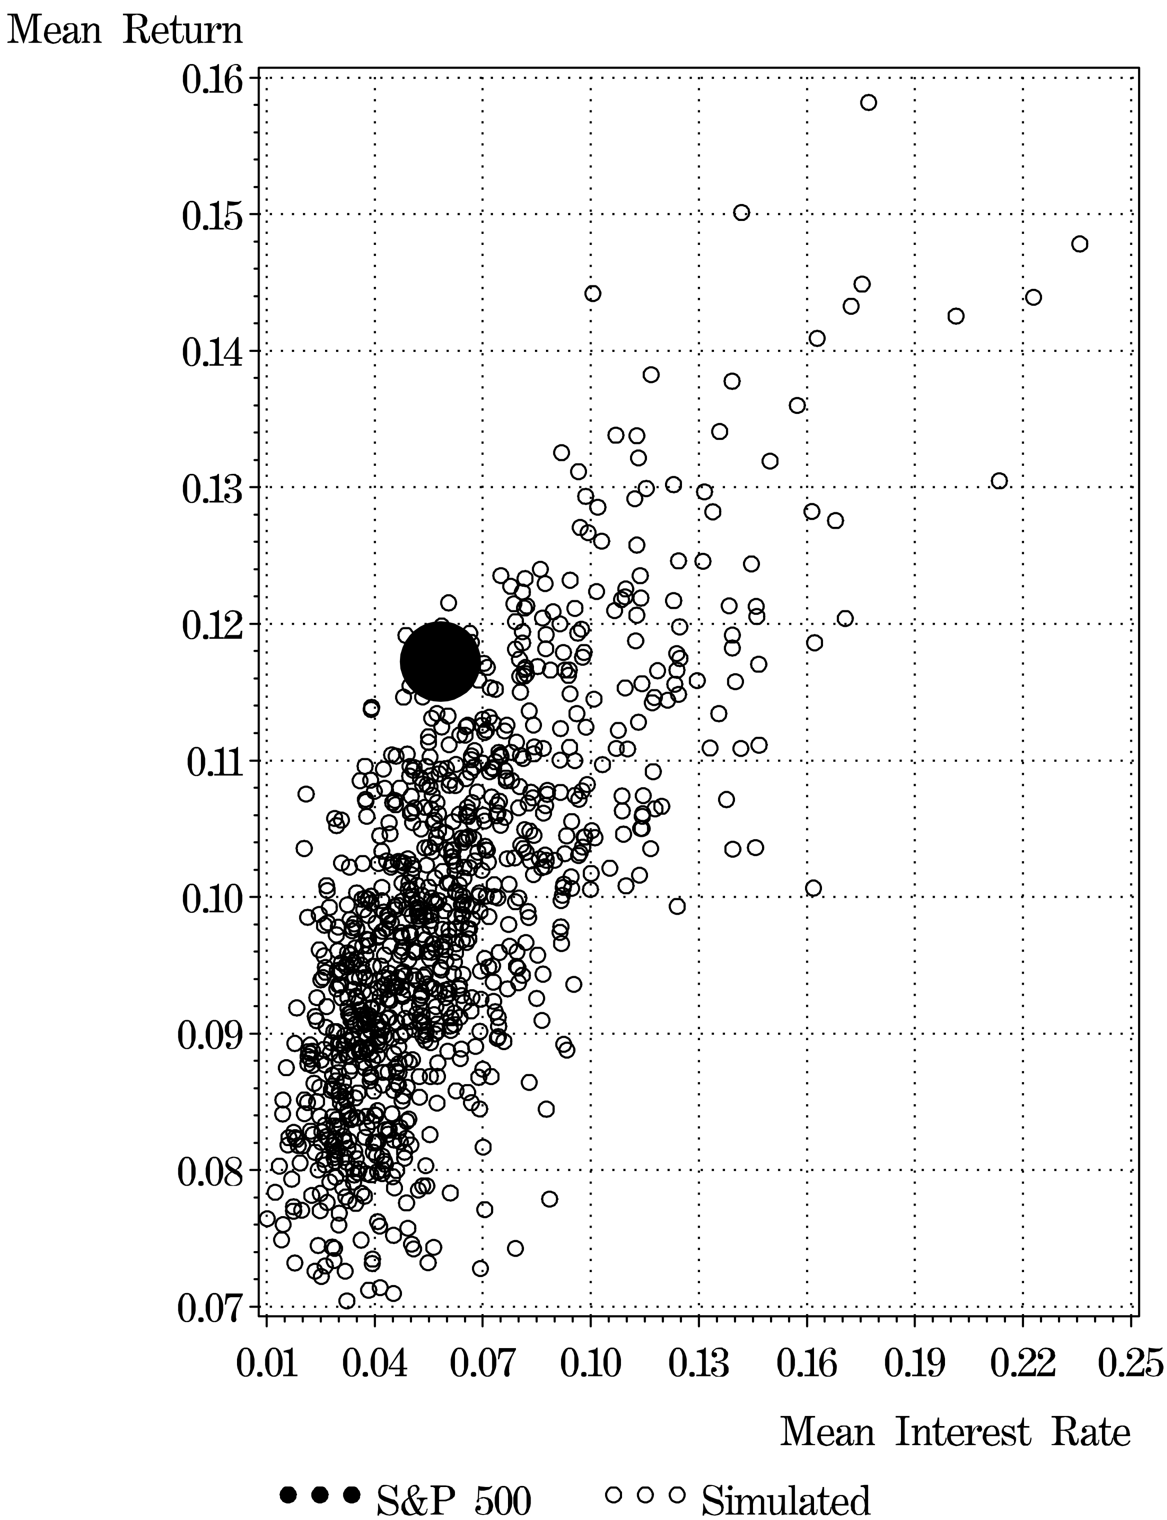

Figure 4: Return Mean versus Interest Rate Mean

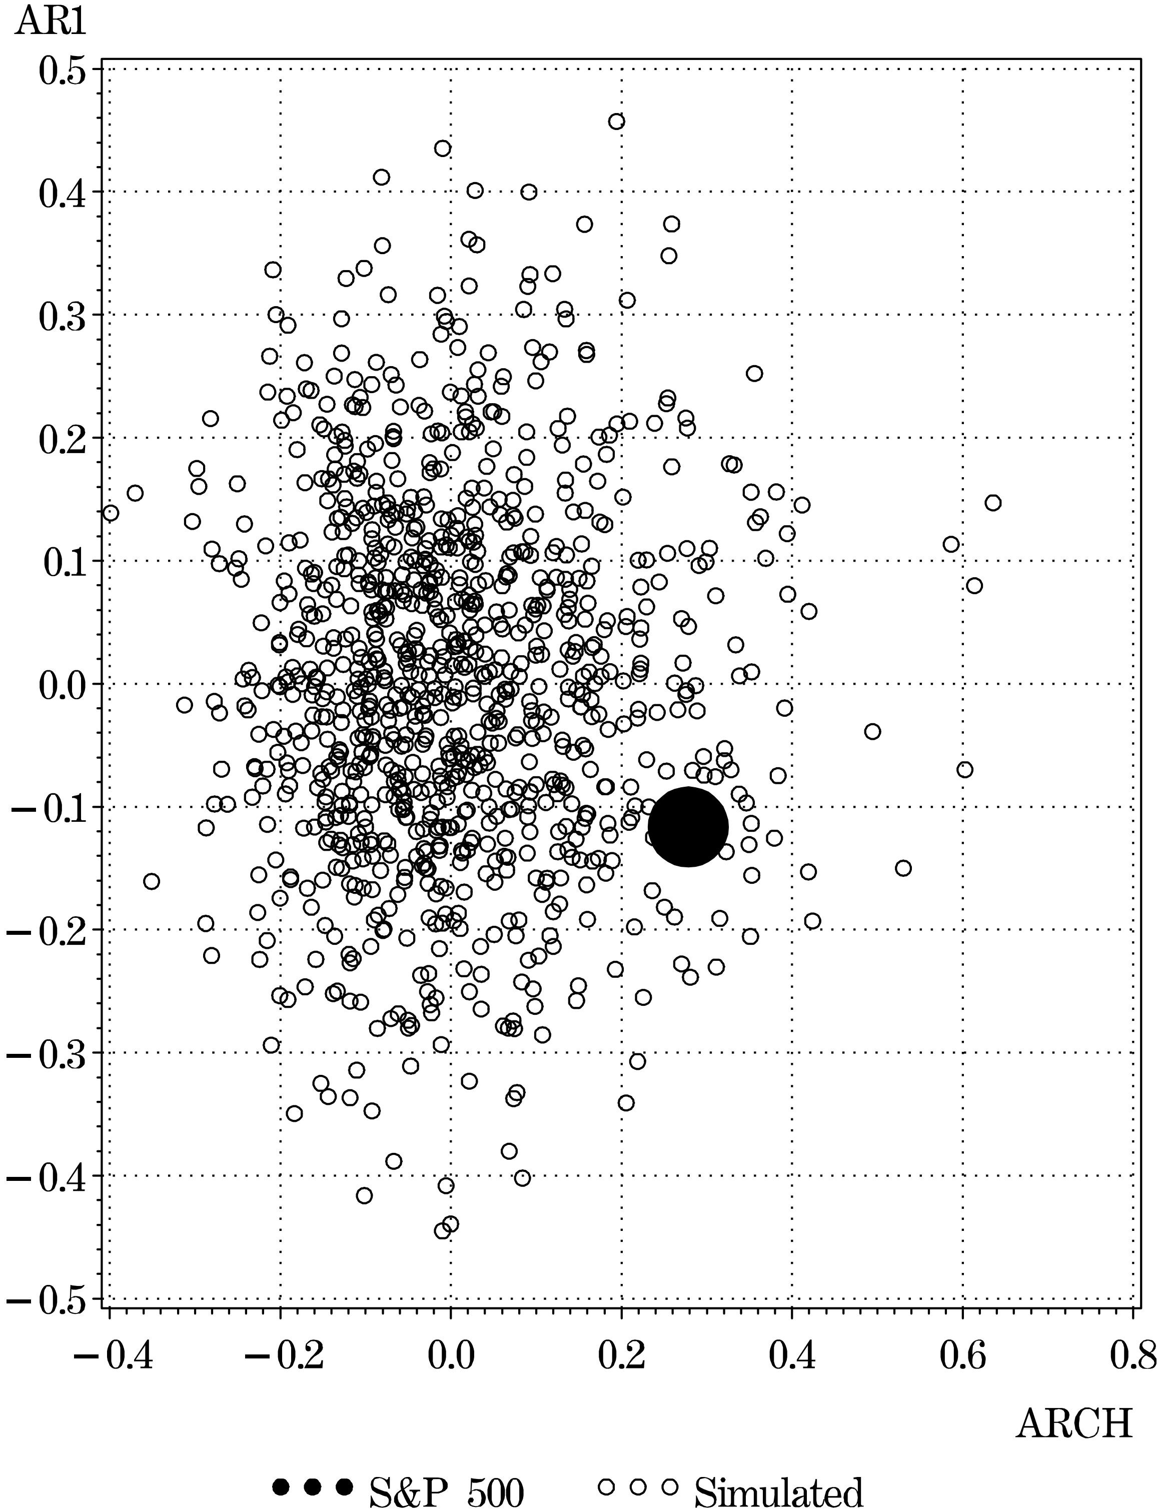

Figure 5: Return AR(1) Coefficient versus ARCH(1) Coefficient