46. Nelson CC, Kim MJ (1993) Predictable stock returns: The role of small sample bias. J Finance 43:641–661
47. Pastor L, Stambaugh RF (2001) The equity premium and structural breaks. J Finance 56(4):1207–1239
48. Pastor L, Stambaugh RF (2006) Predictive systems: Living with imperfect predictors, graduate School of Business. University of Chicago Journal of Finance (forthcoming)
49. Paye BS, Timmermann A (2006) Instability of return prediction models. J Empir Finance 13(3):274–315
50. Piazzesi M, Schneider M, Tuzel S (2007) Housing, consumption, and asset pricing. J Financial Econ 83(March):531–569
51. Polk C, Thompson S, Vuolteenaho T (2006) Cross-sectional forecasts of the equity risk premium. J Financial Econ 81: 101–141
52. Poterba JM, Summers LH (1988) Mean reversion in stock returns: Evidence and implications. J Financial Econ 22:27–60
53. Rozeff MS (1984) Dividend yields are equity risk premia. J Portfolio Manag 49:141–160
54. Shiller RJ (1984) Stock prices and social dynamics. Brook Pap Econ Act 2:457–498
55. Stambaugh RF (1999) Predictive regressions. J Financial Econ 54:375–421
56. Summers LH (1986) Does the stock market rationally reflect fundamental values? J Finance 41:591–601
57. Torous W, Volkanov R, Yan S (2004) On predicting returns with nearly integrated explanatory variables. J Bus 77:937–966
58. Valkanov R (2003) Long-horizon regressions: Theoretical results and applications. J Financial Econ 68:201–232
59. van Binsbergen J, Koijen RS (2007) Predictive regressions: A present-value approach. Duke University (working paper)
60. Viceira L (1996) Testing for structural change in the predictability of asset returns. Harvard University (unpublished manuscript)
61. Vuolteenaho T (2000) Understanding the aggregate book-market ratio and its implications to current equity-premium expectations. Harvard University (unpublished paper)

Books and Reviews

Campbell JY, Lo AW, MacKinlay C (1997) The Econometrics of Financial Markets. Princeton University Press, Princeton

Cochrane JH (2005) Asset Pricing. Princeton University Press, Princeton, NJ

Malkiel BG (2004) A Random Walk Down Wall Street. W.W. Norton, New York

Financial Economics, Time Variation in the Market Return

Mark J. Kamstra[1], Lisa A. Kramer[2]
[1] Schulich School of Business, York University, Toronto, Canada
[2] Rotman School of Management, University of Toronto, Toronto, Canada

Article Outline

Glossary

AR(k) An autoregressive process of order k; a time series model allowing for first order dependence; for instance, an AR(1) model is written as $y_t = \alpha + \rho_1 y_{t-1} + \epsilon_t$ where α and ρ are parameters, ρ is typically assumed to be less than 1 in absolute value, and ϵ_t is an innovation term, often assumed to be Gaussian, independent, and identically distributed over t.

ARCH(q) A special case of the GARCH(p, q) model (see below) where $p = 0$.

Basis point A hundredth of one percent.

Bootstrap A computer intensive resampling procedure, where random draws with replacement from an original sample are used, for instance to perform inference.

Discount rate The rate of return used to discount future cashflows, typically calculated as a risk-free rate (e. g. the 90-day US T-bill rate) plus an equity risk premium.

Equity premium puzzle The empirical observation that the ex post equity premium (see entry below) is higher than is indicated by financial theory.

Ex ante equity premium The extra return investors *expect* they will receive for holding risky assets, over and above the return they would receive for holding a risk-free asset like a Treasury bill. "Ex ante" refers to the fact that the expectation is formed in advance.

Ex post equity premium The extra return investors received *after* having held a risky asset for some period of time. The ex post equity premium often differs from the ex ante equity premium due to random events that impact a risky asset's return.

Free cash flows Cash flows that could be withdrawn from a firm without lowering the firm's current rate of growth. Free cash flows are substantially different from accounting earnings and even accounting measures of the cash flow of a firm.

Fundamental valuation The practice of determining a stock's intrinsic value by discounting cash flows to their present value using the required rate of return.

GARCH(p, q) Generalized autoregressive conditional heteroskedasticity of order (p, q), where p is the order of the lagged variance terms and q is the order of the lagged squared error terms; a time series model allowing for dependence in the conditional variance of a random variable, y. A GARCH(1,1) model is specified as:

$$y_t = \alpha + \epsilon_t; \quad \epsilon_t \sim \left(0, h_t^2\right)$$
$$h_t^2 = \theta + \beta h_{t-1}^2 + \gamma \epsilon_{t-1}^2 ,$$

where α, θ, β, and γ are parameters and ϵ_t is an innovation term.

Market anomalies Empirical regularities in financial market prices or returns that are difficult to reconcile with conventional theories and/or valuation methods.

Markov model A model of a probabilistic process where the random variable can only take on a finite number of different values, typically called states.

Method of moments A technique for estimating parameters (like parameters of the conditional mean and conditional variance) by matching sample moments, then solving the equations for the parameters to be estimated.

SAD Seasonal Affective Disorder, a medical condition by which reduced daylight in the fall and winter leads to seasonal depression for roughly ten percent of the world's population.

Sensation seeking A measure used by psychologists to capture an individual's degree of risk tolerance. High sensation-seeking tendency correlates with low risk tolerance, including tolerance for risk of a financial nature.

Simulated method of moments A modified version of the method of moments (see entry above) that is based on Monte Carlo simulation, used in situations when the computation of analytic solutions is infeasible.

Definition of the Subject

The realized return to any given asset varies over time, occasionally in a dramatic fashion. The value of an asset, its *expected* return, and its volatility, are of great interest to investors and to policy makers. An asset's expected return in excess of the return to a riskless asset (such as a short-term US Treasury bill) is termed the equity premium. The value of the equity premium is central to the valuation of risky assets, and hence a much effort has been devoted to determining the value of the equity premium, whether it varies, and if it varies, how predictable it is. Any evidence of predictable returns is either evidence of a predictably varying equity premium (say, because risk varies predictably) or a challenge to the rationality of markets and the efficient allocation of our society's scarce resources.

In this article, we start by considering the topic of valuation, with emphasis on simulation-based techniques. We consider the valuation of income-generating assets in the context of a constant equity premium, and we also explore the consequences of allowing some time-variation and predictability in the equity premium. Next we consider the equity premium puzzle, discussing a simulation-based technique which allows for precise estimation of the value of the equity premium, and which suggests some constraints on the types of models that should be used for specifying the equity premium process. Finally, we focus on evidence of seasonally varying expected returns in financial markets. We consider evidence that as a whole either presents some challenges to traditional hypotheses of efficient markets, or suggests agents' risk tolerance may vary over time.

Introduction

The pricing of a firm is conceptually straightforward. One approach to valuing a firm is to use historical dividend payments and discount rate data to forecast future payments and discount rates. Restrictions on the dividend and discount rate processes are typically imposed to produce an analytic solution to the fundamental valuation equation (an equation that involves calculating the expectation of an infinite sum of discounted dividends).

Common among many of the available valuation techniques is some form of consideration of multiple scenarios, including good and bad growth and discount rate evolutions, with valuation based on a weighted average of prices from the various scenarios. The valuation technique we focus some attention on, the Donaldson and Kamstra [14] (henceforth DK) methodology, is similar to pricing path-dependent options, as it utilizes Monte Carlo simulation techniques and numerical integration of the possible paths followed by the joint processes of dividend growth and discount rates, explicitly allowing path-dependence of the evolutions. The DK method is very similar in spirit to other approaches in the valuation literature which consider multiple scenarios. One distinguishing feature of the DK methodology we consider is the technique it employs for modeling the discount rate.

Cochrane [9] highlights three interesting approaches for modeling the discount rate: a constant discount rate, a consumption-based discount rate, and a discount rate equal to some variable reference return plus a risk premium. Virtually the entire valuation literature limits its attention to the constant discount rate case, as constant discount rates lead to closed-form solutions to many valuation formulas. DK explore all three methods for modeling the discount rate and find they lead to qualitatively similar results. However, their quantitative results indicate an overall better fit to the price and return data when using a reference return plus a risk premium. Given DK's findings, we use a discount rate equal to some variable reference return plus a risk premium. In implementing this approach for modeling the discount rate used in valuation, it is simplest to assume a *constant* equity premium is added to the reference rate, in particular since the reference rate is permitted to vary (since it is typically proxied using a variable rate like the three-month US T-bill rate). We do not, however, restrict ourselves to the constant equity premium case.

Using the simulation-based valuation methodology of DK and the method of simulated moments, we explore the evidence for a time-varying equity premium and its implications for a long-standing puzzle in financial economics, the equity premium puzzle of Mehra and Prescott [51]. Over the past century the average annual return to investing in the US stock market has been roughly 6% higher than the return to investing in risk-free US T-bills. Making use of consumption-based asset-pricing models, Mehra and Prescott argue that consumption within the US has not been sufficiently volatile to warrant such a large premium on risky stocks relative to riskless bonds, leading them to describe this large premium as the "equity premium puzzle."

Utilizing simulations of the distribution from which ex post equity premia are drawn, conditional on various possible values for investors' ex ante equity premium and calibrated to S&P 500 dividends and US interest rates, we present statistical tests that show a true ex ante equity premium as low as 2% could easily produce ex post premia of 6%. This result is consistent with the well-known observation that ex post equity premia are observed with error, and a large range of realized equity premia are consistent with any given value of the ex ante equity premium. Examining the marginal and joint distributions of financial statistics like price-dividend ratios and return volatility that arise in the simulations versus actual realizations from the US economy, we argue that the range of ex ante equity premia most consistent with the US market data is very close to 3.5%, and the ex ante equity premium process is very unlikely to be constant over time.

A natural question to ask is why might the equity premium fluctuate over time? There are only two likely explanations: changing risk or changing risk aversion. Evidence from the asset-pricing literature, including [20,37,49], and many others shows that priced risk varies over time. We explore some evidence that risk aversion itself may vary over time, as revealed in what is often termed market anomalies. Market anomalies are variations in expected returns which appear to be incongruous with variations in discount rates or risk. The most stark anomalies have to do with deterministic asset return seasonalities, including seasonalities at the weekly frequency such as the weekend effect (below-average equity returns on Mondays), annual effects like the above-average equity returns typically witnessed in the month of January, and other effects like the lower-than-average equity returns often witnessed following daylight saving time-change weekends, and opposing cyclicality in bond versus equity returns correlated to the length of day (known as the SAD effect). We briefly review some of these outstanding puzzles, focusing our attention on the SAD effect and the daylight saving effect.

Valuation

Overview

We begin our discussion of valuation with a broad survey of the literature, including dividend-based valuation, relative valuation, and accounting-based methods. We introduce dividend-based valuation first.

Fundamental valuation techniques that utilize dividends in a *discrete* time framework include Gordon [25], Hawkins [30], Michaud and Davis [53], Farrell [22], Sorensen and Williamson [73], Rappaport [63], Barsky and DeLong [2], Hurley and Johnson [33], [34], Donaldson and Kamstra [14], and Yao [78]. Invariably these approaches are partial equilibrium solutions to the valuation exercise. Papers that use *continuous* time tools to evaluate the fundamental present value equation include Campbell and Kyle [6], Chiang, Davidson, and Okuney [8], Dong and Hirshleifer [17], and Bakshi and Chen [3]. The Dong and Hirshleifer [17] and Bakshi and Chen [3] papers conduct valuation by assuming dividends are proportional to earnings and then modeling earnings. Continuous time papers in this literature typically start with the representative agent/complete markets economic paradigm. Models are derived from primitive assumptions on markets and preferences, such as the equilibrium condition that there exist no arbitrage opportunities, dividend (cash flow) growth rates follow an Ornstein–Uhlenbeck mean-reverting process, and preferences over consumption are represented by the log utility function. Time-varying stochastic discount rates (i. e. the pricing kernel) fall out of the marginal rate of utility of consumption in these models, and the solution to the fundamental valuation problem is derived with the same tools used to price financial derivatives. A critique of dividend-discounting methods is that dividends are typically smoothed and are set low enough so that the dividend payments can be maintained through economic downturns. Authors such as Hackel and Livnat (see p. 9 in [27]) argue that these sorts of considerations imply that historical records of dividend payments may therefore be poor indicators of future cash payments to investors.

A distinct valuation approach, popular amongst practitioners, determines the value of inactively traded firms by finding an actively traded firm that has similar risk, profitability, and investment-opportunity characteristics and then multiplying the actively traded firm's price-earnings (P/E) ratio by the inactively traded firm's earnings. This approach to valuation is often referred to as the relative value method or the constant P/E model. References to this sort of approach can be found in textbooks like [4], and journal articles such as [60,62].

There are also several valuation approaches that are based on the book value of equity, abnormal earnings, and free-cash flows. These approaches are linked to dividends and hence to formal fundamental valuation by well-established accounting relationships. They produce price estimates by valuing firm assets and income streams. The most popular of this class of techniques include the residual income and free-cash-flow methods. See [23,57,61] for further information. All of these valuation methods implicitly or explicitly take the present value of the stream of firm-issued dividends to the investor. The motivation for considering accounting relationships is that these accounting measures are not easily manipulated by firms and so should reflect more accurately the ability of firms to generate cashflows and hence allow more accurate assessments of the fundamental value of a firm than techniques based on dividends.

Fundamental Valuation Methods in Detail

Now that we have surveyed the valuation literature in general, we turn to a formal derivation of several *fundamental* valuation techniques. Investor rationality requires that the current market price P_t of a stock which will pay a per share dividend (cash payment) D_{t+1} one period from now and then sell for P_{t+1}, discounting payments received during period t (i. e., from the beginning of period t to the beginning of period $t+1$) at rate r_t, must satisfy Eq. (1):

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

$\mathcal{E}_t$ is the expectations operator conditional on information available up to the end of period t. Solving Eq. (1) forward under the transversality condition that the expected present value of P_{t+k} goes to zero as k goes to infinity (a "no-bubble" assumption) produces the familiar result that the market price equals the expected present value of future dividends (cash payments); i. e.,

$$P_t = \sum_{k=0}^{\infty} \mathcal{E}_t \left\{ \left(\prod_{i=0}^{k} \left[\frac{1}{1 + r_{t+i}} \right] \right) D_{t+k+1} \right\} . \tag{2}$$

Defining the growth rate of dividends from the beginning of period t to the beginning of period $t+1$ as $g_t^{\mathrm{d}} \equiv (D_{t+1} - D_t)/D_t$ it follows that

$$P_t = D_t \mathcal{E}_t \left\{ \sum_{k=1}^{\infty} \left(\prod_{i=0}^{k} \left[\frac{1 + g_{t+i}^{\mathrm{d}}}{1 + r_{t+i}} \right] \right) \right\} . \tag{3}$$

Equation (3) is the fundamental valuation equation, which is not controversial and can be derived under the law of one price and non-satiation alone, as by Rubinstein [69] and others. Notice that the cash payments D_{t+k} include all cash disbursements from the firm, including cash dividends and share repurchases. Fundamental valuation methods based directly on Eq. (3) are typically called dividend discount models.

Perhaps the most famous valuation estimate based on Eq. (3) comes from the Gordon [25] Growth Model. If dividend growth rates and discount rates are constant, then it is straightforward to derive the Gordon fundamental price estimate from Eq. (3):

$$P_t^{\mathrm{G}} = D_t \left[\frac{1 + g^{\mathrm{d}}}{r - g^{\mathrm{d}}} \right] , \tag{4}$$

where r is the constant discount rate value and g^{d} is the (conditionally) constant growth rate of dividends. To produce the Gordon Growth Model valuation estimate, all we need are estimates of the dividend growth rate and discount rate, which can be obtained in a variety of ways, including the use of historically observed dividends and returns.

Extensions of the Gordon Growth Model exploit the fundamental valuation equation, imposing less stringent assumptions. The simple Gordon Growth Model imposes a constant growth rate on dividends (dividends are expected to grow at the same rate every period) while Hurley and Johnson [33] and [34] and Yao [78] develop Markov models (models that presume a fixed probability of, say, maintaining the dividend payment at current levels, and a probability of raising it, thus incorporating more realistic dividend growth processes). Two examples of these models found in Yao [78] are the Additive Markov Gordon model (Eq. (1) of Yao [78] and the Geometric Markov Gordon model (Eq. (2) of Yao [78]). These models can be interpreted as considering different scenarios for dividend growth for a particular asset, estimating the appropriate price for the asset under each scenario, and then averaging the prices using as weights the probability of given scenarios being observed.

The Additive Markov Gordon Growth Model is:

$$P_t^{\mathrm{ADD}} = D_t/r + \left[1/r + (1/r)^2\right]\left(q^{\mathrm{u}} - q^{\mathrm{d}}\right)\Delta\,, \tag{5}$$

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

The Geometric Markov Gordon Growth Model is:

$$P_t^{\mathrm{GEO}} = D_t\left[\frac{1+(q^{\mathrm{u}}-q^{\mathrm{d}})\Delta^{\%}}{r-(q^{\mathrm{u}}-q^{\mathrm{d}})\Delta^{\%}}\right], \tag{6}$$

where $\Delta^{\%} = \sum_{t=2}^{T}|(D_t - D_{t-1})/D_{t-1}|/(T-1)$ is the average absolute value of the percentage rate of change in the dividend payment.

The method of DK is also an extension of the Gordon Growth Model, taking the discounted dividend growth model of Eq. (3) and re-writing it as

$$P_t = D_t \sum_{k=0}^{\infty} \mathcal{E}_t\left\{\prod_{i=0}^{k} y_{t+i}\right\}, \tag{7}$$

where $y_{t+i} = (1+g_{t+i}^{\mathrm{d}})/(1+r_{t+i})$ is the discounted dividend growth rate. Under the DK method, the fundamental price is calculated by forecasting the range of possible evolutions of y_{t+i} up to some distant point in the future, period $t+I$, calculating $PV = D_t\sum_{k=0}^{I}(\prod_{i=0}^{k} y_{t+i})$ for each possible evolution of y_{t+i}, and averaging these values of PV across all the possible evolutions. (The value of I is chosen to produce a very small truncation error. Values of $I = 400$ to 500 for annual data have been found by DK to suffice). In this way, the DK approach mirrors other extensions of the Gordon Growth Model. It is primarily distinguished from other approaches that extend the Gordon Growth Model in two regards. First, more sophisticated time series models, estimated with historical data, are used to generate the different outcomes (scenarios) by application of Monte Carlo simulation. Second, in contrast to typical modeling in which only dividend growth rates vary, the joint evolution of cashflow growth rates and discount rates are explicitly modeled as time-varying.

Among the attractive features of the free-cash-flow and residual income valuation methods is that they avoid the problem of forecasting dividends, by exploiting relationships between accounting data and dividends. It is the practical problem of forecasting dividends to infinity that have led many researchers to explore methods based on accounting data. See, for instance, Penman and Sougiannis [61].

Assume a flat term structure (i. e., a constant discount rate $r_t = r$ for all t) and write

$$P_t = \sum_{k=1}^{\infty} \frac{\mathcal{E}_t\{D_{t+k}\}}{(1+r)^k}\,. \tag{8}$$

The clean-surplus relationship relating dividends to earnings is invoked in order to derive the residual income model:

$$B_{t+k} = B_{t+k-1} + E_{t+k} - D_{t+k}\,, \tag{9}$$

where B_{t+k} is book value and E_{t+k} is earnings per share. Solving for D_{t+k} in Eq. (9) and substituting into Eq. (8) yields

$$P_t = \sum_{k=1}^{\infty} \frac{\mathcal{E}_t\{B_{t+k-1} + E_{t+k} - B_{t+k}\}}{(1+r)^k}\,,$$

or

$$\begin{aligned} P_t &= B_t + \sum_{k=1}^{\infty} \frac{\mathcal{E}_t\{E_{t+k} - r\cdot B_{t+k-1}\}}{(1+r)^k} - \frac{\mathcal{E}_t\{B_{t+\infty}\}}{(1+r)^{\infty}} \\ &= B_t + \sum_{k=1}^{\infty} \frac{\mathcal{E}_t\{E_{t+k} - r\cdot B_{t+k-1}\}}{(1+r)^k}\,, \end{aligned} \tag{10}$$

where $B_{t+\infty}/(1+r)^{\infty}$ is assumed to equal zero. $E_{t+k} - r\cdot B_{t+k-1}$ is termed abnormal earnings.

To derive the free cash flow valuation model, we relate dividends to cash flows with a financial assets relation in place of the clean surplus relation:

$$fa_{t+k} = fa_{t+k-1} + i_{t+k} + c_{t+k} - D_{t+k}\,, \tag{11}$$

where fa_{t+k} is financial assets net of financial obligations, i_{t+k} is interest revenues net of interest expenses, and c_{t+k} is cash flows realized from operating activities net of investments in operating activities, all of which can be positive or negative. A net interest relation is often assumed,

$$i_{t+k} = r fa_{t+k-1}\,. \tag{12}$$

See Fetham and Ohlson [23] for further discussion. Solving for D_{t+k} in Eq. (11) and substituting into Eq. (8), utilizing Eq. (12) and assuming the discounted present value of financial assets fa_{t+k} goes to zero as k increases, yields the free-cash-flow valuation equation:

$$P_t = fa_t + \sum_{k=1}^{\infty} \frac{\mathcal{E}_t\{c_{t+k}\}}{(1+r)^k}\,. \tag{13}$$

More on the Fundamental Valuation Method of Donaldson and Kamstra

A number of approaches can be taken to conduct valuation using the DK model shown in Eq. (7). By imposing a very simple structure for the conditional expectation of discounted dividend growth rate (y_t in Eq. (7)), the expression can be solved analytically, for instance by assuming that the discounted dividend growth rate is a constant. As shown by DK, however, analytic solutions become complex for even simple ARMA models, and with sufficient non-linearity, the analytics can be intractable. For this reason, we present a general solution algorithm based on the DK method of Monte Carlo simulation.

This method simulates y_t into the future and performs a numerical (Monte Carlo) integration to estimate the terms $\{\prod_{k=0}^{i} y_{t+k}\}$ where $y_{t+k} = (1 + g^{\mathrm{d}}_{t+k})/(1 + r_{t+k})$ in the classic case of a dividend-paying firm. A general heuristic is as follows:

Step I: Model y_t, $t = 1, \ldots, T$, as conditionally time-varying, for instance as an AR(k)-GARCH(p, q) process, and use the estimated model to make conditional mean forecasts $\hat{y}_t$, $t = 1, \ldots, T$, and variance forecasts, conditional on data observed only before period t. Ensure that this model is consistent with theory, for instance that the mean level of y is less than one. This mean value can be calibrated to available data, such as the mean annual y value of 0.94 observed in the last 50 years of S&P 500 data. Recall that although analytic solutions are available for simple processes, the algorithm presented here is applicable to virtually arbitrarily non-linear conditional processes for the discounted cash payment rate y.

Step IIa: Simulate discounted cash payment growth rates. That is, produce ys that might be observed in period t given what is known at period $t-1$. To do this for a given period t, simulate a population of J independent possible shocks (say draws from a normal distribution with mean zero and appropriate variance, or bootstrapped from the data) $\epsilon_{t,j}$, $j = 1, \ldots, J$, and add these shocks separately to the conditional mean forecast $\hat{y}_t$ from Step I, producing $y_{t,j} = \hat{y}_t + \epsilon_{t,j}$, $j = 1, \ldots, J$. The result is a simulated cross-section of J possible realizations of y_t standing at time $t-1$, i.e. different paths the economy may take next period.

Step IIb: Use the estimated model from Step I to make the conditional mean forecast $\hat{y}_{t+1,j}$, conditional on only the jth realization for period t (i.e., $y_{t,j}$ and $\epsilon_{t,j}$) and the data known at period $t-1$, to form $y_{t+1,j}$.

Step IIc: Repeat Step IIb to form $y_{t+2,j}, y_{t+3,j}, \ldots, y_{t+I,j}$ for each of the J economies, where I is the number of periods into the future at which the simulation is truncated. Form the perfect foresight present value ($P^*_{t,j}$) for each of the J possible economies:

$$P^*_{t,j} = A_t\Big(y_{t,j} + y_{t,j}y_{t+1,j} + y_{t,j}y_{t+1,j}y_{t+2,j} + \cdots + \prod_{i=0}^{I} y_{t+i,j}\Big)\,; \quad j = 1, \ldots, J\,.$$

Provided I is chosen to be large enough, the truncated terms $\prod_{i=0}^{K} y_{t+i,j}$, $K = I+1, \ldots, \infty$ will be negligible.

Step III: Calculate the DK fundamental price for each $t = 1, \ldots, T$:

$$P^{\mathrm{DK}}_t = \sum_{j=1}^{J} P^*_{t,j}/J\,. \tag{14}$$

These fundamental price estimates P^{DK}_t can be compared to the actual price (if market prices exist) at the beginning of period t to test for bubbles as demonstrated by DK, or if period t is the future, P^{DK}_t is the fundamental price forecast. This procedure is represented diagrammatically in Exhibit 1.

To illustrate the sort of forecasts that can be produced using this technique, we illustrate graphically the S&P 500 index over the past 100 years together with predicted values based on the Gordon Growth Model and the DK method. The free-cash-flow and residual income methods are not easily adapted to forecasting index prices like the S&P 500, and so are omitted here. The type of data depicted in the following figure is described in some detail by Kamstra [39].

Figure 1 has four panels. In the panels, we plot the level of the S&P 500 index (marked with bullets and a solid line) alongside price forecasts from each of the valuation techniques. In Panel A we plot the index together with the basic Gordon Growth Model price forecasts (marked with stars), in Panels B and C we plot the index together with the Additive and Geometric Gordon Growth Models' forecasts (with squares and triangles respectively), and in Panel D we plot the index alongside the DK method's forecasts (marked with diamonds). In each panel the price scale is logarithmic.

We see in Panels A, B, and C that the use of the any of the Gordon models for forming annual forecasts of the S&P 500 index level produces excessively smooth price forecasts. (If we had plotted return volatility, then the market returns would appear excessively volatile in comparison to to forecasted returns). Evidence of periods of inflated market prices relative to the forecasted prices, i.e.,

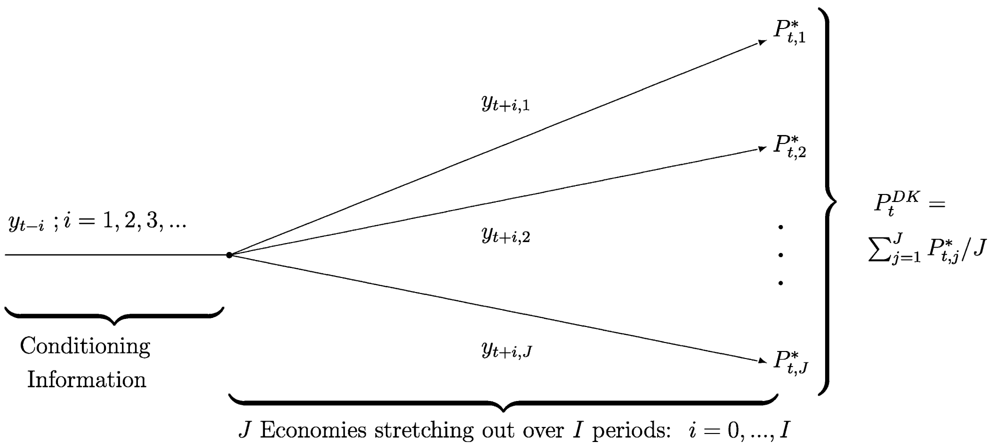

Financial Economics, Time Variation in the Market Return, Exhibit 1
Diagram of DK Monte Carlo integration

evidence of price bubbles, is apparent in the periods covering the 1920s, the 1960s, the last half of the 1980s, and the 1990s. However, if the Gordon models are too simple (since each Gordon-based model ignores the forecastable nature of discount rates and dividend growth rates), then this evidence may be misleading.

In Panel D, we see that the DK model is better able to capture the volatility of the market, including the boom of the 1920s, the 1960s and the 1980s. The relatively better performance of the DK price estimate highlights the importance of accounting for the slow fade rate of dividend growth rates and discount rates, i. e., the autocorrelation of these series. The failure of the DK method to capture the height of the 1990s boom leaves evidence of surprisingly high prices during the late 1990s. If the equity premium fell in the 1990s, as some researchers have speculated (see for instance Pástor and Stambaugh [59]), then all four sets of the plotted fundamental valuation forecasts would be expected to produce forecasts that undershoot actual prices in the 1990s, as all these methods incorporate a constant equity premium. If this premium were set too high, future cashflows would be discounted too aggressively, biasing the valuation methods downward.

The Equity Premium Puzzle

The fact that all four fundamental valuation methods we consider spectacularly fail to capture the price boom of the 1990s, possibly as a result of not allowing a time-varying equity premium, sets the stage to investigate the equity premium puzzle of Mehra and Prescott [51]. The equity premium is the extra return, or premium, that investors demand in order to be compelled to purchase risky stock instead of risk-free debt. We call this premium the ex ante equity premium (denoted π_e), and it is formally defined as the difference between the expected return on risky assets, $\mathcal{E}\{R\}$, and the expected risk-free rate, $\mathcal{E}\{r_f\}$:

$$\pi_e \equiv \mathcal{E}\{R\} - \mathcal{E}\{r_f\} \ . \tag{15}$$

The ex post equity premium is typically estimated using historical equity returns and risk-free rates, as we do not observe the ex ante premium. 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. A standard approach to calculate ex post equity premium, $\hat{\pi}_e$, is:

$$\hat{\pi}_e \equiv \overline{R} - \overline{r}_f \, . \tag{16}$$

Of course it is unlikely that the stock return we estimate ex post equals investors' anticipated ex ante return. Thus a 6% ex post equity premium in the US data may not be a challenge to economic theory. The question we ask is therefore: if investors' true ex ante premium is X%, what is the probability that the US economy could randomly produce an ex post premium of at least 6%? We can then argue whether or not the 6% ex post premium observed in the US data is consistent with various ex ante premium values, X%, with which standard economic theory may be more compatible. We can also consider key financial statistics and yields from the US economy to investigate if an X% ex ante equity premium could likely be consistent with the combinations that have been observed, such as high Sharpe ratio and low dividend yields, low interest rates and high ex post equity premia, and so on.

Authors have investigated the extent to which ex ante considerations may impact the realized equity premium.

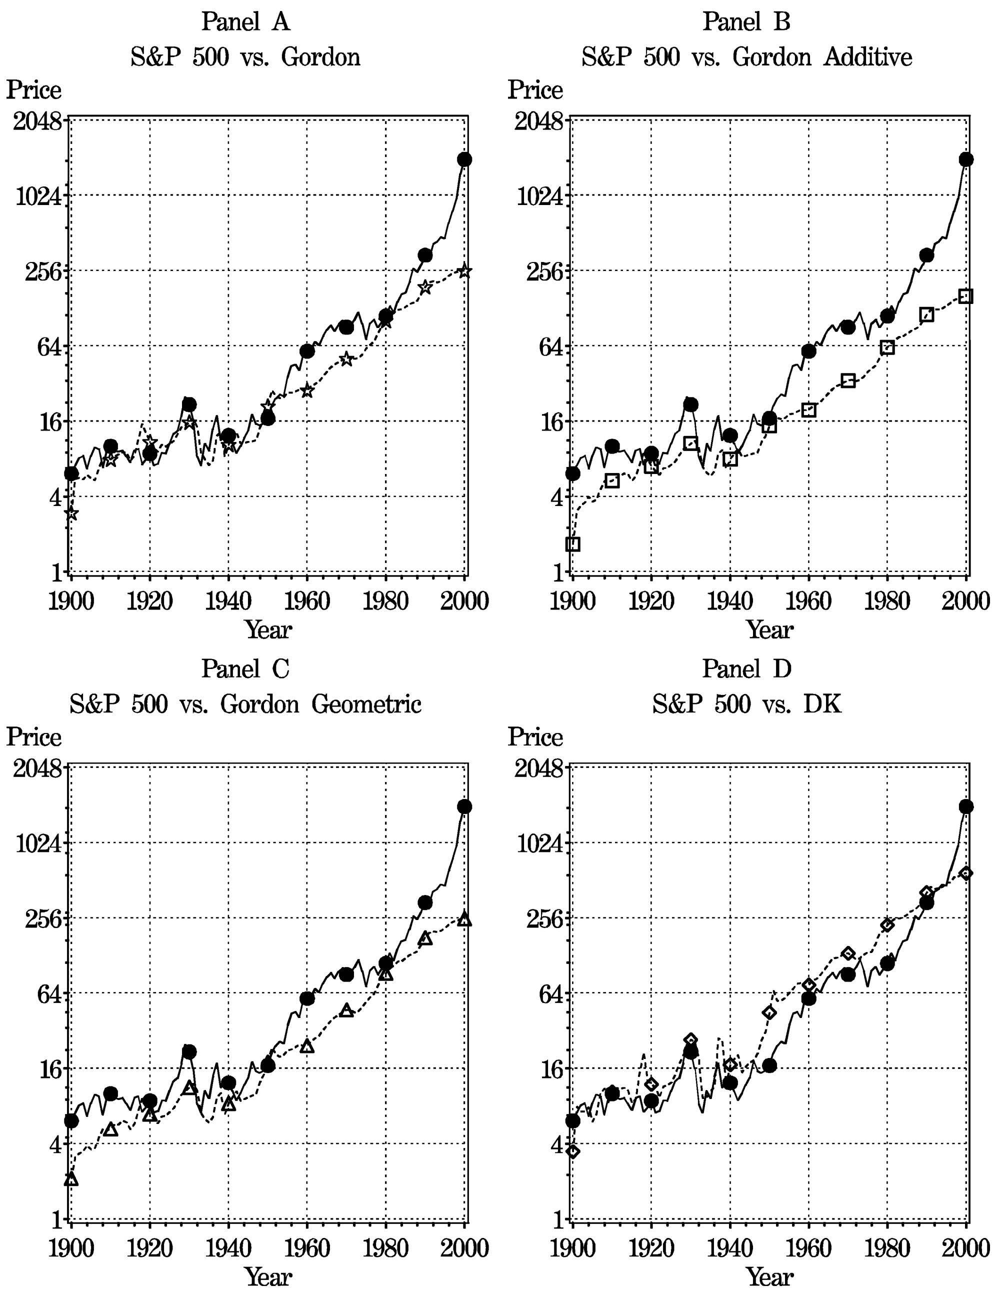

Financial Economics, Time Variation in the Market Return, Figure 1
S&P 500 index level versus price forecasts from four models. S&P 500 index: ●, Gordon Growth price: ★, Additive Gordon Growth price: □, Geometric Gordon Growth price: △, DK price ◇

For example, Rietz [65] investigated the effect that the fear of a serious, but never realized, depression would have on equilibrium asset prices and equity premia. Jorion and Goetzmann [38] take the approach of comparing the US stock market's performance with stock market experiences in many other countries. They find that, while some markets such as the US and Canada have done very well over the past century, other countries have not been so fortunate; average stock market returns from 1921 to 1996 in France, Belgium, and Italy, for example, are all close to zero, while countries such as Spain, Greece, and Romania have experienced negative returns. It is difficult, however, to conduct statistical tests because, first, the stock indices Jorion and Goetzmann consider are largely contemporaneous and returns from the various indices are not independent. Statistical tests would have to take into account the panel nature of the data and explicitly model covariances across countries. Second, many countries in the comparison pool are difficult to compare directly to the United States in terms of economic history and underlying data generating processes. (Economies like Egypt and Romania, for example may have equity premia generated from data generating processes that differ substantially from that of the US).

There are some recent papers that make use of fundamental information in examining the equity premium. One such paper, Fama and French [21], uses historical dividend yields and other fundamental information to calculate estimates of the equity premium which are smaller than previous estimates. Fama and French obtain point estimates of the ex post equity premium ranging from 2.55% (based on dividend growth rate fundamentals) to 4.78% (based on bias-adjusted earnings growth rate fundamentals), however these estimates have large standard errors. For example, for their point estimate of 4.32% based on non-bias-adjusted earnings growth rates, a 99% confidence interval stretches from approximately −1% to about 9%. Mehra and Prescott's [51] initially troubling estimate of 6% is easily within this confidence interval and is in fact within one standard deviation of the Fama and French point estimate.

Calibrating to economy-wide dividends and discount rates, Donaldson, Kamstra, and Kramer [16] employ simulation methods similar to DK to simulate a distribution of possible price and return outcomes. Comparing these simulated distributions with moments of the actual data then permits them to test various models for the equity premium process. Could a realized equity premium of 6% be consistent with an ex ante equity premium of 2%? Could an ex ante equity premium of 2% have produced the low dividend yields, high ex post equity premia, and high Sharpe ratios observed in the US over the last half century?

A summary of the basic methodology implemented by Donaldson, Kamstra, and Kramer [16], is as follows:

(a) Assume a mean value for the equity premium that investors demand when they first purchase stock (e. g., 2%) and a time series process for the premium, say a deterministic drift downward in the premium of 5 basis points per year, asymptoting no lower than perhaps 1%. This assumed premium is added to the risk-free interest rate to determine the discount rate that an investor would rationally apply to a forecasted dividend stream in order to calculate the present value of dividend-paying stock.
(b) Estimate econometric models for the time-series processes driving dividends and interest rates in the US economy (and, if necessary, for the equity premium process), allowing for autocorrelation and covariation. Then use these models to Monte Carlo simulate a variety of potential paths for US dividends, interest rates, and equity premia. The simulated 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.
(c) Given the assumed process for the equity premium investors demand ex ante (which is the same for all simulated economies in a given experiment), use a discounted-dividend model to calculate the fundamental stock returns (and hence ex post equity premia) that arise in each simulated economy. All economies have the same ex ante equity premium process, and yet all economies 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, it is feasible to calculate various other important characteristics, like Sharpe ratios and dividend yields.
(d) Examine the distribution of ex post equity premia, interest rates, dividend growth rates, Sharpe ratios, and dividend yields that arise conditional on various values of the ex ante equity premia. Comparing the performance of the US economy with intersections of the various univariate and multivariate distributions of these quantities and conducting joint hypothesis tests allows the determination of a narrow range of equity premia consistent with the US market data. Note that this is the method of simulated moments, which is well adapted to estimate the ex ante equity premium. The simulated method of moments was developed by McFadden [50] and Pakes and Pollard [58]. Duffie and Singleton [18] and Corradi and Swanson [11] employ simulated method of moments in an asset pricing context.

Further details on the simulation methodology are provided by Donaldson, Kamstra, and Kramer [16]. They make use of annual US stock and Treasury data observed from 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. The model that generated the data we use to illustrate this simulation methodology is Model 6 of Donaldson, Kamstra, and Kramer [16], a model that allows for trending, autocorrelated, and co-varying dividend growth rates, interest rates and equity premia, as well as for a structural break in the equity premium process. We show later that allowing for trends and structural breaks in the equity premium process is a crucial factor in the model's ability to capture the behavior of the observed US market data.

We focus on the intuition behind the Donaldson, Kamstra, and Kramer technique by looking at bivariate plots of simulated data, conditional on various values of the ex ante equity premium. In every case, the pair of statistics we plot are dependent on each other in some way, allowing us to make interesting conditional statements. Among the bivariate distributions we consider, we will see some that serve primarily to confirm the ability of our simulations to produce the character and diversity of results observed in US markets. Some sets of figures rule out ex ante equity premia below 2.5% while others rule out ex ante equity premia above 4.5%. Viewed collectively, the figures serve to confirm that the range of ex ante equity premia consistent with US market data is in the close vicinity of 3.5%.

Figure 2 contains joint distributions of mean returns and return standard deviations arising in our simulations based on four particular values of the ex ante equity premium (2.5% in Panel A, 3.5% in Panel B, 4.5% in Panel C, and 6% in Panel D). Each panel contains a scatter plot of two thousand points, with each point representing a pair of statistics (mean return and return standard deviation) arising in one of the simulated half-century economies. The combination based on the US realization is shown in each plot with a crosshair (a pair of solid straight lines with the intersection marked by a solid dot). The set of simulated pairs in each panel is surrounded by an ellipse which represents a 95% bivariate confidence bound, based on the asymptotic normality (or log-normality, where appropriate) of the plotted variables. (The 95% confidence ellipsoids are asymptotic approximations based on joint normality of the sample estimates of the moments of the simulated data. All of the sample moment estimates we consider are asymptotically normally distributed, as can be seen by appealing to the appropriate law of large numbers). The confidence ellipse for the 2.5% case is marked with diamonds, the 3.5% case with circles, the 4.5% case with squares, and the 6% case with circled crosses.

Notice that the sample mean for the US economy (the intersection of the crosshairs) lies loosely within cloud of points that depict the set of simulated economies for each ex ante equity premium case. That is, our simulations produce mean returns and return volatility that roughly match the US observed moments of returns, *without our having calibrated to returns*. Notice also that the intersection of the crosshairs is outside (or very nearly outside) the 95% confidence ellipse in all cases except that of the 3.5% ex ante equity premium. (In unreported results that study a finer grid of ex ante equity premium values, we found that only those simulations based on values of the ex ante equity premium between about 2.5% and 4.5% lead to 95% confidence ellipses that encompass the US economy crosshairs. As the value of the ex ante equity premium falls below 2.5% or rises above 4.5%, the confidence ellipse drifts further away from the crosshairs). Based on this set of plots, we can conclude that ex ante equity premia much less than or much greater than 3.5% are inconsistent at the 5% confidence level with the observed mean return and return volatility of S&P 500 returns. χ^2 tests presented in Donaldson, Kamstra, and Kramer [16] confirm this result.

We can easily condense the information contained in these four individual plots into one plot, as shown in Panel A of Fig. 3. The scatterplot of points representing individual simulations are omitted in the condensed plot, but the confidence ellipses themselves (and the symbols used to distinguish between them) are retained. Panel A of Fig. 3 repeats the ellipses shown in Fig. 2, so that again we see that only the 3.5% ex ante equity premium case is well within the confidence ellipse at the 5% significance level. In presenting results for additional bivariate combinations, we follow the same practice, omitting the points that represent individual simulations and using the same set of symbols to distinguish between confidence ellipses based on ex ante equity premia of 2.5%, 3.5%, 4.5%, and 6%.

In Panel B of Fig. 3 we consider the four sets of confidence ellipses for mean return and mean dividend yield combinations. Notice that as we increase the ex ante equity premium, the confidence ellipses shift upward and to the right. Notice also that with higher values of the ex ante equity premium we tend to have more variable dividend yields. That is, the confidence ellipse covers a larger range of dividend yields when the value of the ex ante equity premium is larger. The observed combination of S&P 500 mean return and mean dividend yield, represented by the intersecting crosshairs, lies within the confidence ellipse for the 2.5% and 3.5% cases, very close to the ellipse for the 4.5% case, and far outside the ellipse for the 6% case.

Panel C of Fig. 3 plots confidence ellipses for mean interest rates versus mean ex post equity premia. The intersection of the crosshairs is within all four of the shown confidence ellipses. As we calibrate our model to the US interest rate, and as the ex post equity premium has a large

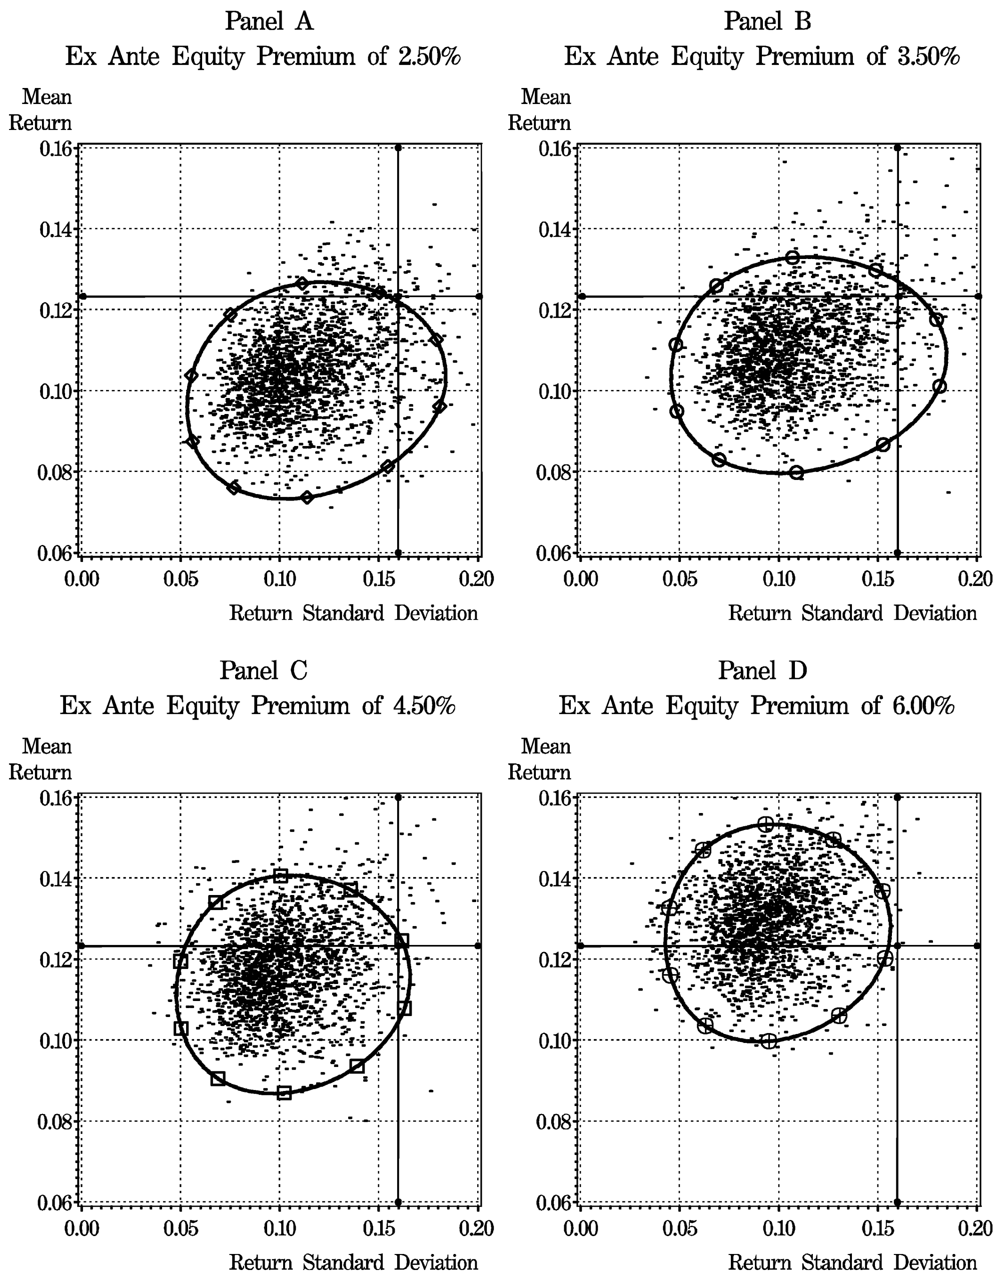

Financial Economics, Time Variation in the Market Return, Figure 2

Bivariate scatterplots of simulated data for a model allowing for trends and structural breaks. The model upon which these scatterplots are based allows for trends and structural breaks in the equity premium process, as well as autocorrelated and co-varying dividend growth rates, interest rates, and equity premia. Observed market data are indicated with crosshairs, and confidence ellipses are marked as follows. Ex ante equity premium of 2.5%: ⋄, Ex ante equity premium of 3.5%: ∘, Ex ante equity premium of 4.5%: □, Ex ante equity premium of 6%: ⊕

variance, it is not surprising that the US experience is consistent with the simulated data from the entire range of ex ante equity premia considered here. This result is merely telling us that the ex post equity premium is not, by itself, particularly helpful in narrowing the possible range for the ex ante equity premium (consistent with the empirical imprecision in measuring the ex post equity premium which has been extensively documented in the literature). Notice as well that the confidence ellipses in Panel C are all negatively sloped: we see high mean interest rates with low eq-

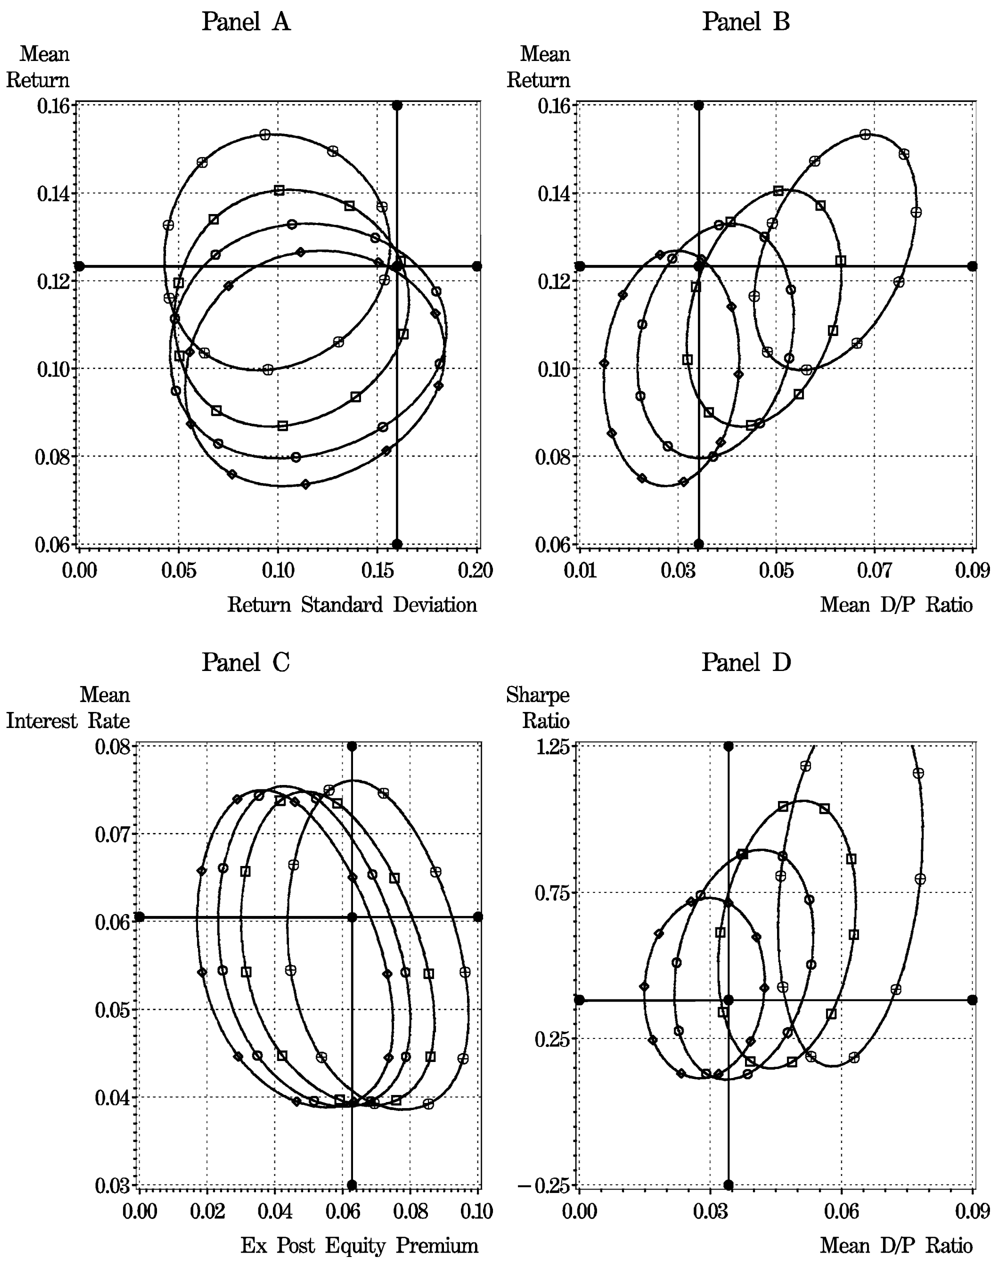

Financial Economics, Time Variation in the Market Return, Figure 3

Confidence ellipses based on simulated data for a model allowing for trends and structural breaks. The model upon which these scatterplots are based allows for trends and structural breaks in the equity premium process, as well as autocorrelated and co-varying dividend growth rates, interest rates, and equity premia. Observed market data are indicated with crosshairs, and confidence ellipses are marked as follows. 2.5% ex post equity premium: ⋄, 3.5% ex post equity premium: ∘, 4.5% ex post equity premium: □, 6% ex post equity premium: ⊕

uity premia and low mean interest rates with high equity premia. Many researchers, including Weil [74], have commented that the flip side of the high equity premium puzzle is the low risk-free rate puzzle. Here we confirm that the dual puzzle arises in our simulated economies as well. It appears that this puzzle is a mechanical artifact coming out of the calculation of the premium. As the ex post equity premium equals the mean return minus the mean interest rate, a decrease in the interest rate, all else held constant, must lead to a higher ex post equity equity premium.

Panel D of Fig. 3 contains the confidence ellipses for the Sharpe ratio (or reward-to-risk ratio, calculated as the average annual difference between the arithmetic return and the risk-free rate divided by the standard deviation of the annual differences) and the mean dividend yield. As the ex ante equity premium is increased from 2.5%, the confidence ellipses shift from being centered on the crosshairs to far to the right of the crosshairs. The US expe-

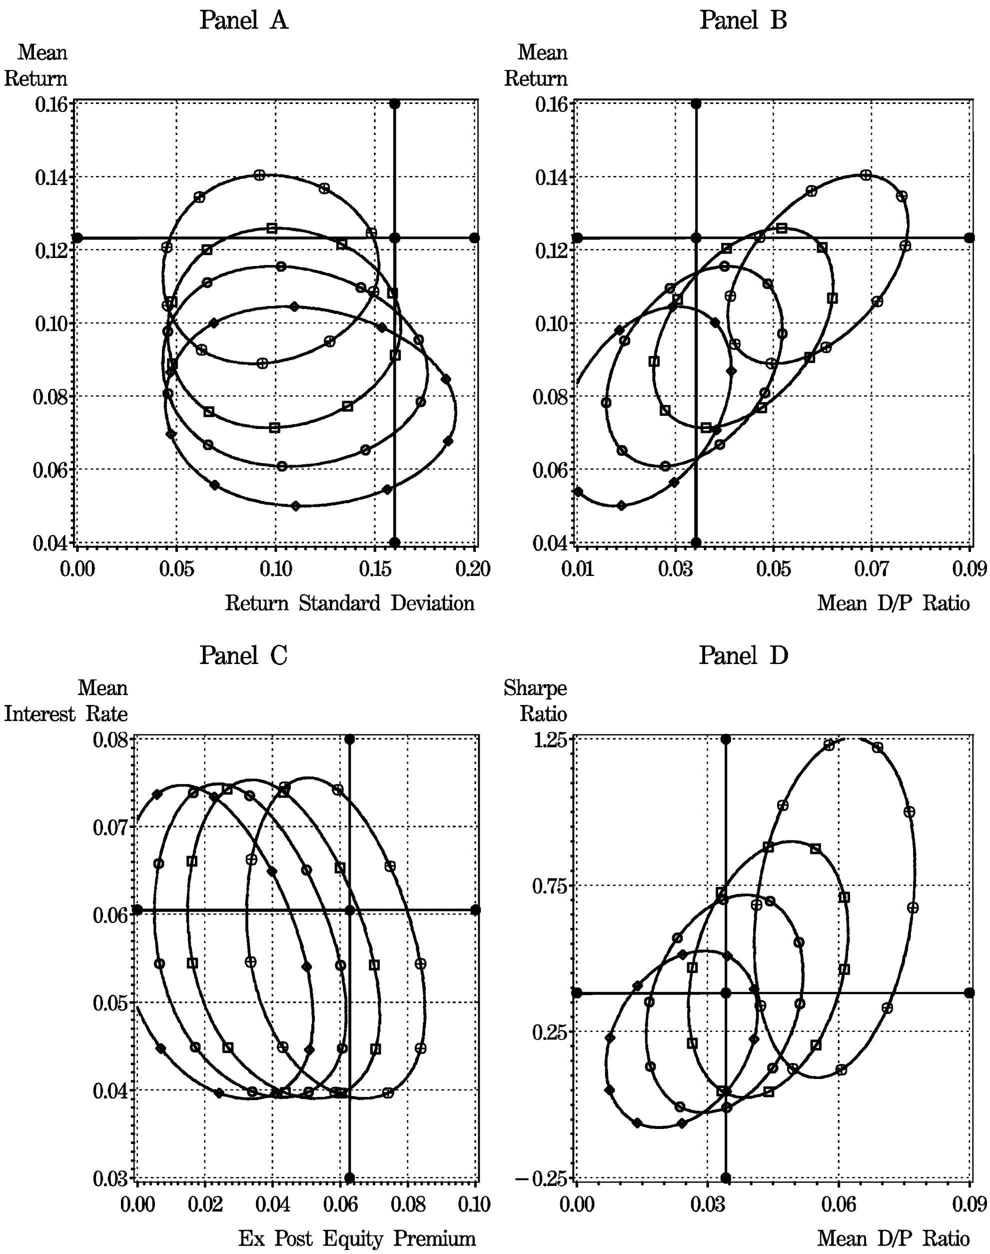

Financial Economics, Time Variation in the Market Return, Figure 4

Confidence ellipses based on simulated data for a restricted model that does not allow for trends and structural breaks. The model upon which these scatterplots are based does not allow for trends or structural breaks in the equity premium process, but does allow for autocorrelated and co-varying dividend growth rates, interest rates, and equity premia. Observed market data are indicated with crosshairs, and confidence ellipses are marked as follows. 2.5% ex post equity premium: ⋄, 3.5% ex post equity premium: ∘, 4.5% ex post equity premium: □, 6% ex post equity premium: ⊕

rience, indicated by the crosshairs at a Sharpe ratio of approximately 0.4 and a mean dividend yield of about 3.5%, is well outside the 95% confidence ellipse for the 6% ex ante equity premium case, suggesting a 6% ex ante equity premium is inconsistent with the jointly observed S&P 500 Sharpe ratio and mean dividend yield. Indeed Fama and French [21] and Jagannathan, McGrattan, and Scherbina [35] make reference to dividend yields to argue that the equity premium may be much smaller than 6%; our analysis gives us a glimpse of just how much smaller it might be.

Overall in Fig. 3, the joint realization of key characteristics of the US market data suggests that the true ex ante equity premium is no lower than 2.5%, no higher than 4.5%, and is most likely near 3.5%. Multivariate χ^2 tests performed by Donaldson, Kamstra, and Kramer [16] indicate a 95% confidence interval of plus-or-minus 50 basis points around 3.5%.

Consider now Fig. 4, which details simulated data from a restricted model that has a time-varying equity premium but no trends or structural breaks. Donaldson, Kamstra, and Kramer [16] study this simplified model and find that it performs poorly relative to the model we consider in Figs. 2 and 3 in terms of its ability to capture the behavior of US market data. Figure 4 shows that with the restricted model, no values of the ex ante equity premium are consistent with the observed US mean return, standard deviation, and dividend yield. That is, the simulation-based mean return and dividend yield ellipses do not contain the US data crosshairs for any value of the ex ante equity premium considered. (χ^2 tests presented in Donaldson, Kamstra, and Kramer [16] strongly support this conclusion). The implication is that it is essential to model trends and structural breaks in the equity premium process in order to accurately capture the dynamics of observed US data. Donaldson, Kamstra, and Kramer show that model failure becomes even more stark if the equity premium is constrained to be constant.

Overall, the evidence in Figs. 3 and 4 does not itself resolve the equity premium puzzle, but evidence in Fig. 3 (based on the model that allows for trends and structural breaks in the equity premium process) does provide a narrow target range of plausible equity premia that economic models should be able to explain. Additionally, the evidence in Figs. 3 and 4 points to a secondary issue ignored in the literature prior to the work of Donaldson, Kamstra, and Kramer [16], that it is crucial to model the equity premium as both time-varying and as having trends and structural breaks. We saw in Fig. 4 that high return volatility, high ex post equity premia, and low dividend yields cannot be explained easily by constant equity premium models. This result has clear implications for valuation: simple techniques that restrict the discount rate to a constant are remarkably inconsistent with the US experience of time-varying equity premia, and serious attention should be paid to modeling a time-varying rate for use in discounting future expected cash flows.

Time-Varying Equity Premia: Possible Biological Origins

To the extent that the simulation techniques considered in the previous section suggest that the equity premium varies over time, it is interesting to consider some empirical evidence of time-varying equity premia. We first survey some examples of high-frequency variations in the equity premium, and then we explore in detail two examples which may arise due to reasons that relate to human biology and/or psychology.

There is a wide range of evidence of high-frequency movement in the equity premium. At the highest frequency, we observe roughly 'U-shaped' intra-day returns (see [29,36,77]), with returns being perhaps somewhat higher during the morning trading period than in the afternoon (see [46]). At the weekly frequency, returns from Friday's close until Monday's close are low and even negative on average, as first identified by Cross [12]. Rogalski [66] found prices rose during Mondays, thus identifying the negative average realizations that followed Fridays as a weekend effect and not a Monday effect. Turning to the monthly domain, Ogden [56] documented a turn of the month effect where returns in the first half of the month are higher than returns in the second half of the month. At the annual frequency, there is the well-known turn-of-the-year effect, first shown by Rozeff and Kinney [68]. Keim [45] showed that half of the year's excess returns for small firms arose in January, and half of the January returns took place in the first five days of the month. Further, Reinganum [64] showed that January returns are higher for small firms whose price performed poorly in the previous year. All of this is consistent with the tax-loss-selling hypothesis whereby investors realize losses at the end of the tax year, leading to higher returns in January after the tax-loss selling ends.

Next we turn our attention to two cases of time-varying equity premia that may arise for reasons related to human physiology. One is Seasonal Affective Disorder (SAD), and another is daylight saving time changes.

Seasonal Affective Disorder

Past research suggests there are seasonal patterns in the equity premium which may arise due to cyclical changes in the risk tolerance of individual investors over the course of the year related to SAD. The medical condition of SAD, according to Rosenthal [67], is a recurrent depression associated with diminished daylight in the fall, affecting many millions of Americans, as well as peoples from around the world, even those located near the equator. (In a study of 303 patients attending a primary care facility in Vancouver, Schlager, Froom, and Jaffe [70] found that 9% were clinically diagnosed with SAD and another 29% had significant winter depressive symptoms without meeting conditions for major depression. Other studies have found similar magnitudes, though some research has found that prevalence varies by latitude, with more extreme latitudes having a larger proportion of SAD-sufferers.) SAD is classified as a major depressive disorder. The symptoms of SAD include anxiety, periods of sadness, chronic fatigue, difficulty concentrating, lethargy, sleep disturbance, sugar and carbohydrate craving and associated weight gain, loss of interest in sex, and of course, clinical depression. Psychologists have shown that depressed people have less tolerance for risk in general. (See [7,32,82,83]). Psychologists refer to risk tolerance in terms of "sensation seeking" tendency, measured using a scale developed by Zuckerman [80], [81]. Those who tolerate (or seek) high levels of risk tend to score high on the sensation-seeking scale. Differences in sensation-seeking tendencies have been linked to gender (see [5] for example), race (see [31] for instance), age (see, for example, [84]), and other personal characteristics.

Economists and psychologists working together have shown that sensation-seeking tendency translates into tolerance for risk of a specifically financial or economic nature. For instance, Wong and Carducci [76] find that individuals who score low on tests of sensation seeking display greater risk aversion in making financial decisions, including the decision to purchase stocks, bonds, and insurance. Harlow and Brown [28] document the link between sensation seeking and financial risk tolerance by building on results from psychiatry which show that high blood levels of a particular enzyme are associated with depression and a lack of sensation seeking while low levels of the enzyme are associated with a high degree of sensation seeking. Harlow and Brown write, "Individuals with neurochemical activity characterized by lower levels of [the enzyme] and with a higher degree of sensation-seeking are *more willing to accept economic risk* … Conversely, high levels of this enzyme and a low level of sensation seeking appear to be associated with risk-averse behavior." (pp. 50–51, emphasis added). These findings suggest an individual's level of sensation seeking is indicative of his or her tolerance for financial risk.

Given these relationships, Kamstra, Kramer, and Levi [42] conjecture that during the fall and winter seasons, when a fraction of the population suffers from SAD, the proportion of risk-averse investors rises. Risk-averse investors shun risky stocks in the fall, they argue, which has a negative influence on stock prices and returns. As winter progresses and daylight becomes more plentiful, investors start to recover from their depression and become more willing to hold risky assets, at which time stock prices and returns should be positively influenced.

If the extent or severity of SAD is greater at more extreme latitudes, then the SAD effect on stock returns should be greater in stock markets at high latitudes and less in markets close to the equator. Also, the pattern of returns in the Southern Hemisphere should be the opposite of that in the Northern Hemisphere as are the seasons. Thus, Kamstra, Kramer and Levi [42] study stock market indices for the US, Sweden, Britain, Germany, Canada, New Zealand, Japan, Australia, and South Africa. They regress each country's daily stock returns on a variety of standard control variables plus two variables intended to capture the impact of SAD on returns. The first of these two variables, SAD_t, is a simple function of the length of night at the latitude of the respective market for the fall and winter months for which SAD has been documented to be most severe. The second of these variables, a fall dummy variable denoted Fall_t, is included because the SAD hypothesis implies the expected effect on returns is different before versus after winter solstice. Specifically, when agents initially become more risk averse, they should shun risky assets which should cause prices to be lower than would otherwise be observed, and when agents revert to normal as daylight becomes more plentiful, prices should rebound. The result should be lower returns in the autumn, higher returns in the winter, and thus a high equity premium for investors who hold through the autumn and winter periods. The Fall_t dummy variable is used to capture the lower autumn returns. Both SAD_t and Fall_t are appropriately defined for the Southern Hemisphere countries, accounting for the six month difference in seasons relative to the Northern Hemisphere markets.

Table 1 summarizes the average annual effect due to each of the SAD variables, SAD_t and Fall_t, for each of the international indices Kamstra, Kramer, and Levi [42] study. For comparison, the unconditional average annual return for each index is also provided. Observe that the annualized return due to SAD_t is positive in every country, varying from 5.7 to 17.5 percent. The SAD effect is generally larger the further are the markets from the equator. The negative annualized returns due to Fall_t demonstrate the fact that SAD typically causes returns to be

Financial Economics, Time Variation in the Market Return, Table 1
Average annual percentage return due to SAD variables

Country (Index)	Annual return due to SAD_t	Annual return due to $fall_t$	Unconditional annual return
US (S&P 500)	9.2***	−3.6**	6.3***
Sweden (Veckans Affärar)	13.5**	−6.9**	17.1***
Britain (FTSE 100)	10.3**	−2.3	9.6***
Germany (DAX 30)	8.2*	−4.3**	6.5**
Canada (TSX 300)	13.2***	−4.3**	6.1***
New Zealand (Capital 40)	10.5**	−6.6**	3.3
Japan (NIKKEI 225)	6.9*	−3.7**	9.7***
Australia (All ordinaries)	5.7	0.5	8.8***
South Africa (Datastream global index)	17.5*	−2.1	14.6***

One, two, and three asterisks denote significantly different from zero at the ten, five, and one percent level respectively, based on one-sided tests. Source: Table 3 in [42].

shifted from the fall to the winter. Garrett, Kamstra, and Kramer [24] study seasonally-varying risk aversion in the context of an equilibrium asset pricing model, allowing the price of risk to vary with length of night through the fall and winter seasons. They find the risk premium on equity varies through the seasons in a manner consistent with investors being more risk averse due to SAD in the fall and winter.

Kamstra, Kramer, and Levi [43] show that there is an opposite seasonal pattern in Treasury returns relative to stock returns, consistent with time-varying risk aversion being the underlying force behind the seasonal pattern previously shown to exist in stock returns. If SAD-affected investors are shunning risky stocks in the fall as they become more risk averse, then they should be favoring safe assets at that time, which should lead to an opposite pattern in Treasury returns relative to stock returns. The seasonal cycle in the Treasury market is striking, with a variation of more than 80 basis points between the highest and lowest average monthly returns. The highest Treasury returns are observed when equity returns are lowest, and *vice versa*, which is a previously unknown pattern in Treasury returns.

Kamstra, Kramer, and Levi [43] define a new measure which is linked directly to the clinical incidence of SAD. The new measure uses data on the weekly or monthly onset of and recovery from SAD, obtained from studies of SAD patients in Vancouver and Chicago conducted by medical researchers. Young, Meaden, Fogg, Cherin, and Eastman [79] and Lam [47] document the clinical *onset* of SAD symptoms and *recovery* from SAD symptoms among North Americans known to be affected by SAD. Young et al. study 190 SAD-sufferers in Chicago and find that 74 percent of them are first diagnosed with SAD in the weeks between mid-September and early November. Lam studies 454 SAD patients in Vancouver on a monthly basis and finds, that the peak timing of diagnosis is during the early fall. Lam [47] also studies the timing of clinical remission of SAD and finds it peaks in April, with almost half of all SAD-sufferers first experiencing complete remission in that month. March is the second most common month for subjects to first experience full remission, corresponding to almost 30 percent of subjects. For most SAD patients, the initial onset and full recovery are separated by several months over the fall and winter.

Direct use of Kamstra, Kramer, and Levi's [43] variable (which is an estimate of population-wide SAD onset/recovery based on specific samples of individuals) could impart an error-in-variables problem (see [48]), thus they utilize an instrumented version detailed in the paper, which they call Onset/Recovery, denoted $\hat{OR}_t$. The instrumented SAD measure $\hat{OR}_t$ reflects the change in the proportion of SAD-affected individuals actively suffering from SAD. The measure is defined year-round (unlike the original Kamstra, Kramer, and Levi [42], SAD_t variable, which is defined for only the fall and winter months), taking on positive values in the summer and fall and negative values in the winter and spring. Its value peaks near the fall equinox and reaches a trough near the spring equinox. (The exact monthly values of $\hat{OR}_t$ are reported by Kamstra, Kramer, and Levi [43].) The opposite signs on $\hat{OR}_t$ across the fall and winter seasons should, in principle, permit it to capture the opposite impact on equity or Treasury returns across the seasons, without use of a dummy variable. Kamstra, Kramer, and Levi [43] find that use of $\hat{OR}_t$ as a regressor to explain seasonal patterns in Treasury and equity returns renders the SAD_t and $Fall_t$ (used by Kamstra, Kramer, and Levi [42]) as economically and statistically insignificant, suggesting the Onset/Recovery variable does a far better job of explaining seasonal variation in returns than the original proxies which are not directly related to the incidence of SAD.

Kamstra, Kramer, and Levi [43] show that the seasonal Treasury and equity return patterns are unlikely to arise from macroeconomic seasonalities, seasonal variation in risk, cross-hedging between equity and Treasury markets, investor sentiment, seasonalities in the Treasury market auction schedule, seasonalities in the Treasury debt supply, seasonalities in the Federal Reserve Board's interest-rate-setting cycle, or peculiarities of the sample period considered. They find that the seasonal cycles in equity and Treasury returns become more pronounced during periods of high market volatility, consistent with time-varying risk aversion among market participants. Furthermore, they apply the White [75] reality test and find that the correlation between returns and the clinical incidence of seasonal depression cannot be easily dismissed as the simple result of data snooping.

DeGennaro, Kamstra, and Kramer [13] and Kamstra, Kramer, and Levi [13] provide further corroborating evidence for the hypothesis that SAD leads to time variation in financial markets by considering (respectively) bid-ask spreads for stocks and the flow of funds in and out of risky and safe mutual funds. In both papers they find strong support for the link between seasonal depression and time-varying risk aversion.

Daylight Saving Time Changes

The second potential biological source of time-varying equity premia we consider arises on the two dates of the year when most of the developed world shifts clocks forward or backward an hour in the name of daylight saving. Psychologists have found that changes in sleep patterns (due to shift work, jet lag, or daylight saving time changes, for example) are associated with increased anxiety, which is suggestive of a link between changes in sleep habits and time-varying risk tolerance. See [26,52], and citations found in [10] and [72] for more details on the link between sleep disruptions and anxiety. In addition to causing heightened anxiety, changes in sleep patterns also inhibit rational decision-making, lower one's information-processing ability, affect judgment, slow reaction time, and reduce problem-solving capabilities. Even a change of one hour can significantly affect behavior.

Kamstra, Kramer, and Levi [40] explore the financial market ramifications of a link between daylight saving time-change-induced disruptions in sleep patterns and individuals' tolerance for risk. They find, consistent with psychology studies that show a gain or loss of an hour's sleep leads to increased anxiety, investors seem to shun risky stock on the trading day following a daylight saving time change. They consider stock market indexes from four countries where the time changes happen on non-overlapping dates, the US, Canada, Britain, and Germany. Based on stock market behavior over the past three decades, the authors find that the magnitude of the average return on spring daylight saving weekends is typically between two to five times that of ordinary weekends, and the effect is even stronger in the fall. Kamstra, Kramer, and Levi [41] show that the effect is not driven by a few extremely negative observations, but rather the entire distribution of returns shifts to the left following daylight saving time changes, consistent with anxious investors selling risky stock.

Future Directions

We divide our discussion in this section into three parts, one for each major topic discussed in the article.

Regarding fundamental valuation, a promising future path is to compare estimates emerging from sophisticated valuation methods to market prices, using the comparison to highlight inconsistencies in the modeling assumptions (such as restrictions on the equity premium used by the model, restrictions on the growth rate imposed for expected cash flows, and the implied values of those quantities that can be inferred from market prices). Even if one believes that markets are efficient and investors are rational, there is still much to be learned from calculating fundamentals using models and examining discrepancies relative to observed market prices.

Regarding the simulation techniques for estimating the equity premium, a promising direction for future research is to exploit these tools to forecast the volatility of stock prices. This may lead to new alternatives to existing option-implied volatility calculations and time-series techniques such as ARCH (for an overview of these methods see [15]). Another fruitful future direction would be to apply the simulation techniques to the valuation of individual companies' stock (as opposed to valuing, say, stock market indexes).

Regarding the topic of time-varying equity premia that may arise for biological reasons, a common feature of both of the examples explored in Sect. "Time-Varying Equity Premia: Possible Biological Origins", SAD and daylight-saving-time-change-induced fluctuations in the risk premium, is that in both cases the empirical evidence is based on aggregate financial market data. There is a recent trend in finance toward documenting phenomena at the individual level, using data such as individuals' financial asset holdings and trades in their brokerage accounts. (See [1,54,55] for instance). A natural course forward is to build upon the existing aggregate market support for the prevalence of time-varying risk aversion by testing at the individual level whether risk aversion varies through the course of the year due to seasonal depression and during shorter intervals due to changes in sleep patterns. An additional potentially fruitful direction for future research is to integrate into classical asset pricing models the notion that biological factors might impact asset returns through changes in agents' degree of risk aversion. That is, human traits such as seasonal depression may lead to regularities in financial markets that are not mere anomalies; rather they may be perfectly consistent with rational agents making sensible decisions given their changing tolerance for risk. This new line of research would be similar in spirit to the work of Shefrin [71] who considers the way behavioral biases like overconfidence can be incorporated into the pricing kernel in standard asset pricing models. While the behavioral biases Shefrin considers typically involve humans making errors, the biological factors described here might be considered rational due to their involvement of time-varying risk aversion.

Bibliography

Primary Literature

1. Barber B, Odean T (2001) Boys will be boys: Gender, overconfidence, and common stock investment. Q J Econ 116: 261–292
2. Barsky RB, DeLong JB (1993) Why does the stock market fluctuate? Q J Econ 108:291–311
3. Bakshi G, Chen Z (2005) Stock valuation in dynamic economies. J Financ Market 8:115–151
4. Brealey RA, Myers SC, Allen F (2006) Principles of corporate finance, 8th edn. McGraw-Hill Irwin, New York
5. Byrnes JP, Miller DC, Schafer WD (1999) Gender differences in risk taking: A meta-analysis. Psychol Bull 125:367–383
6. Campbell JY, Kyle AS (1993) Smart money, noise trading and stock price behavior. Rev Econ Stud 60:1–34
7. Carton S, Jouvent R, Bungener C, Widlöcher D (1992) Sensation seeking and depressive mood. Pers Individ Differ 13: 843–849
8. Chiang R, Davidson I, Okuney J (1997) Some theoretical and empirical implications regarding the relationship between earnings, dividends and stock prices. J Bank Financ 21:17–35
9. Cochrane JH (2001) Asset pricing. Princeton University Press, Princeton
10. Coren S (1996) Sleep Thieves. Free Press, New York
11. Corradi V, Swanson NR (2005) Bootstrap specification tests for diffusion processes. J Econom 124:117–148
12. Cross F (1973) The behavior of stock prices on Fridays and Mondays. Financ Anal J 29:67–69
13. DeGennaro R, Kamstra MJ, Kramer LK (2005) Seasonal variation in bid-ask spreads. University of Toronto (unpublished manuscript)
14. Donaldson RG, Kamstra MJ (1996) A new dividend forecasting procedure that rejects bubbles in asset prices. Rev Financ Stud 9:333–383
15. Donaldson RG, Kamstra MJ (2005) Volatility forecasts, trading volume, and the ARCH versus option-implied volatility trade-off. J Financ Res 28:519–538
16. Donaldson RG, Kamstra MJ, Kramer LA (2007) Estimating the ex ante equity premium. University of Toronto Manuscript
17. Dong M, Hirshleifer DA (2005) A generalized earnings-based stock valuation model. Manchester School 73:1–31
18. Duffie D, Singleton KJ (1993) Simulated moments estimation of Markov models of asset prices. Econometrica 61:929–952
19. Engle RF (1982) Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. Econometrica 50:987–1007
20. Fama EF, French KR (1993) Common risk factors in the returns on stocks and bonds. J Financ Econ 33:3–56
21. Fama EF, French KR (2002) The equity premium. J Financ 57:637–659
22. Farrell JL (1985) The dividend discount model: A primer. Financ Anal J 41:16–25
23. Feltham GA, Ohlson JA (1995) Valuation and clean surplus accounting for operating and financial activities. Contemp Account Res 11:689–731
24. Garrett I, Kamstra MJ, Kramer LK (2005) Winter blues and time variation in the price of risk. J Empir Financ 12:291–316
25. Gordon M (1962) The Investment, Financing and Valuation of the Corporation. Irwin, Homewood
26. Gordon NP, Cleary PD, Parker CE, Czeisler CA (1986) The prevalence and health impact of shiftwork. Am J Public Health 76:1225–1228
27. Hackel KS, Livnat J (1996) Cash flow and security analysis. Irwin, Chicago
28. Harlow WV, Brown KC (1990) Understanding and assessing financial risk tolerance: A biological perspective. Financ Anal J 6:50–80
29. Harris L (1986) A transaction data study of weekly and intradaily patterns in stock returns. J Financ Econ 16:99–117
30. Hawkins DF (1977) Toward an old theory of equity valuation. Financ Anal J 33:48–53
31. Hersch J (1996) Smoking, seat belts and other risky consumer decisions: differences by gender and race. Managerial Decis Econ 17:471–481
32. Horvath P, Zuckerman M (1993) Sensation seeking, risk appraisal, and risky behavior. Personal Individ Diff 14:41–52
33. Hurley WJ, Johnson LD (1994) A realistic dividend valuation model. Financ Anal J 50:50–54
34. Hurley WJ, Johnson LD (1998) Generalized Markov dividend discount models. J Portf Manag 24:27–31
35. Jagannathan R, McGrattan ER, Scherbina A (2000) The declining US equity premium. Fed Reserve Bank Minneap Q Rev 24:3–19
36. Jain PC, Joh G (1988) The dependence between hourly prices and trading volume. J Financ Quant Analysis 23:269–283
37. Jegadeesh N, Titman S (1993) Returns to buying winners and selling losers: Implications for stock market efficiency. J Financ 48:65–91

38. Jorion P, Goetzmann WN (1999) Global stock markets in the twentieth century. J Financ 54:953–980
39. Kamstra MJ (2003) Pricing firms on the basis of fundamentals. Fed Reserve Bank Atlanta Econ Rev First Quarter:49–70
40. Kamstra MJ, Kramer LA, Levi MD (2000) Losing sleep at the market: the daylight saving anomaly. Am Econ Rev 90: 1005–1011
41. Kamstra MJ, Kramer LA, Levi MD (2002) Losing sleep at the market: The daylight saving anomaly: Reply. Am Econ Rev 92:1257–1263
42. Kamstra MJ, Kramer LA, Levi MD (2003) Winter blues: A SAD stock market cycle. Am Econ Rev 93:324–343
43. Kamstra MJ, Kramer LA, Levi MD (2007) Opposing Seasonalities in Treasury versus Equity Returns. University of Toronto Manuscript
44. Kamstra MJ, Kramer LA, Levi MD, Wermers R (2008) Seasonal asset allocation: evidence from mutual fund flows. University of Toronto Manuscript
45. Keim DB (1983) Size-related anomalies and stock return seasonality: Further Empirical Evidence. J Financ Econ 12:13–32
46. Kramer LA (2002) Intraday stock returns, time-varying risk premia, and diurnal mood variation. University of Toronto Manuscript
47. Lam RW (1998) Seasonal Affective Disorder: Diagnosis and management. Prim Care Psychiatry 4:63–74
48. Levi MD (1973) Errors in the variables bias in the presence of correctly measured variables. Econometrica 41:985–986
49. Liu W (2006) A liquidity-augmented capital asset pricing model. J Financ Econ 82:631–671
50. McFadden D (1989) A method of simulated moments for estimation of discrete response models without numerical integration. Econometrica 47:995–1026
51. Mehra R, Prescott EC (1985) The equity premium: A puzzle. J Monet Econ 15:145–161
52. Mellinger GD, Balter MB, Uhlenhuth EH (1985) Insomnia and its treatment: prevalence and correlates. Arch Gen Psychiatry 42:225–232
53. Michaud RO, Davis PL (1982) Valuation model bias and the scale structure of dividend discount returns. J Financ 37: 563–576
54. Odean T (1998) Are investors reluctant to realize their losses? J Financ 53:1775–1798
55. Odean T (1999) Do investors trade too much? Am Econ Rev 89:1279–1298
56. Ogden JP (1990) Turn-of-month evaluations of liquid profits and stock returns: A common explanation for the monthly and January effects. J Financ 45:1259–1272
57. Ohlson JA (1995) Earnings, book values, and dividends in equity valuation. Contemp Acc Res 11:661–687
58. Pakes A, Pollard D (1989) Simulation and the asymptotics of optimization estimators. Econometrica 57:1027–1057
59. Pástor Ľ, Stambaugh R (2001) The equity premium and structural breaks. J Financ 56:1207–1239
60. Penman SH (1998) Combining earnings and book value in equity valuation. Contemp Acc Res 15:291–324
61. Penman SH, Sougiannis T (1998) A comparison of dividend, cash flow and earnings approaches to equity valuation. Contemp Acc Res 15:343–383
62. Peters DJ (1991) Using PE/Growth ratios to develop a contrarian approach to growth stocks. J Portf Manag 17:49–51
63. Rappaport A (1986) The affordable dividend approach to equity valuation. Financ Anal J 42:52–58
64. Reinganum MR (1983) The anomalous stock market behavior of small firms in January. J Financ Econ 12:89–104
65. Rietz TA (1988) The equity risk premium: A solution. J Monet Econ 22:117–31
66. Rogalski RJ (1984) New findings regarding day-of-the-week returns: A note. J Financ 35:1603–1614
67. Rosenthal NE (1998) Winter Blues: Seasonal Affective Disorder: What is It and How to Overcome It, 2nd edn. Guilford Press, New York
68. Rozeff MS, Kinney WR (1976) Capital market seasonality: The case of stock returns. J Financ Econ 3:379–402
69. Rubinstein M (1976) The valuation of uncertain income streams and the pricing of options. Bell J Econ 7:407–425
70. Schlager D, Froom J, Jaffe A (1995) Winter depression and functional impairment among ambulatory primary care patients. Compr Psychiatry 36:18–24
71. Shefrin H (2005) A Behavioral Approach to Asset Pricing. Academic Press, Oxford
72. Spira AP, Friedman L, Flint A, Sheikh J (2005) Interaction of sleep disturbances and anxiety in later life: perspectives and recommendations for future research. J Geriatr Psychiatry Neurol 18:109–115
73. Sorensen EH, Williamson DA (1985) Some evidence on the value of dividend discount models. Financ Anal J 41:60–69
74. Weil P (1989) The equity premium puzzle and the risk-free rate puzzle. J Monet Econ 24:401–421
75. White H (2000) A reality check for data snooping. Econometrica 68:1097–1126
76. Wong A, Carducci B (1991) Sensation seeking and financial risk taking in everyday money matters. J Bus Psychol 5: 525–530
77. Wood RA, McInish TH, Ord JK (1985) An investigation of transactions data for NYSE stocks. J Financ 40:723–741
78. Yao Y (1997) A trinomial dividend valuation model. J Portf Manag 23:99–103
79. Young MA, Meaden PM, Fogg LF, Cherin EA, Eastman CI (1997) Which environmental variables are related to the onset of Seasonal Affective Disorder? J Abnorm Psychol 106: 554–562
80. Zuckerman M (1976) Sensation seeking and anxiety, traits and states, as determinants of behavior in novel situations. In: Sarason IG, Spielberger CD (eds) Stress and Anxiety, vol 3. Hemisphere, Washington DC
81. Zuckerman M (1983) Biological Bases of Sensation Seeking, Impulsivity and Anxiety. Lawrence Erlbaum Associates, Hillsdale
82. Zuckerman M (1984) Sensation seeking: A comparative approach to a human trait. Behav Brain Sci 7:413–471
83. Zuckerman M, Buchsbaum MS, Murphy DL (1980) Sensation seeking and its biological correlates. Psychol Bull 88: 187–214
84. Zuckerman M, Eysenck S, Eysenck HJ (1978) Sensation seeking in England and America: Cross-cultural, age, and sex comparisons. J Consult Clin Psychol 46:139–149

Books and Reviews

Dimson E (1988) Stock Market Anomalies. Cambridge University Press, Cambridge

Kocherlakota NR (1996) The equity premium: It's still a puzzle. J Econ Lit 34:42–71
Mehra R (2003) The equity premium: Why is it a puzzle? Financ Anal J 59:54–69
Mehra R, Prescott EC (2003) The equity premium in retrospect. In: Constantinides GM, Harris M, Stulz RM (eds) Handbook of the Economics of Finance: Financial Markets and Asset Pricing, vol 1B. North Holland, Amsterdam, pp 889–938
Penman S (2003) Financial Statement Analysis and Security Valuation, 2nd edn. McGraw-Hill/Irwin, New York
Siegel JJ, Thaler RH (1997) Anomalies: The equity premium puzzle. J Econ Perspect 11:191–200
Thaler RH (2003) The Winner's Curse: Paradoxes and Anomalies of Economic Life. Princeton University Press, Princeton

Financial Forecasting, Non-linear Time Series in

Gloria González-Rivera, Tae-Hwy Lee
Department of Economics, University of California, Riverside, USA

Article Outline

Glossary

Arbitrage pricing theory (APT) the expected return of an asset is a linear function of a set of factors.

Artificial neural network is a nonlinear flexible functional form, connecting inputs to outputs, being capable of approximating a measurable function to any desired level of accuracy provided that sufficient complexity (in terms of number of hidden units) is permitted.

Autoregressive conditional heteroskedasticity (ARCH) the variance of an asset returns is a linear function of the past squared surprises to the asset.

Bagging short for *b*ootstrap *agg*regat*ing*. Bagging is a method of smoothing the predictors' instability by averaging the predictors over bootstrap predictors and thus lowering the sensitivity of the predictors to training samples. A predictor is said to be unstable if perturbing the training sample can cause significant changes in the predictor.

Capital asset pricing model (CAPM) the expected return of an asset is a linear function of the covariance of the asset return with the return of the market portfolio.

Factor model a linear factor model summarizes the dimension of a large system of variables by a set of factors that are linear combinations of the original variables.

Financial forecasting prediction of prices, returns, direction, density or any other characteristic of financial assets such as stocks, bonds, options, interest rates, exchange rates, etc.

Functional coefficient model a model with time-varying and state-dependent coefficients. The number of states can be infinite.

Linearity in mean the process $\{y_t\}$ is linear in mean conditional on X_t if

$$\Pr\left[\mathbb{E}(y_t|X_t) = X_t'\theta^*\right] = 1 \quad \text{for some } \theta^* \in \mathbb{R}^k .$$

Loss (cost) function When a forecast $f_{t,h}$ of a variable Y_{t+h} is made at time t for h periods ahead, the loss (or cost) will arise if a forecast turns out to be different from the actual value. The loss function of the forecast error $e_{t+h} = Y_{t+h} - f_{t,h}$ is denoted as $c_{t+h}(Y_{t+h}, f_{t,h})$, and the function $c_{t+h}(\cdot)$ can change over t and the forecast horizon h.

Markov-switching model features parameters changing in different regimes, but in contrast with the threshold models the change is dictated by a non-observable state variable that is modelled as a hidden Markov chain.

Martingale property tomorrow's asset price is expected to be equal to today's price given some information set

$$\mathbb{E}(p_{t+1}|\mathcal{F}_t) = p_t .$$

Nonparametric regression is a data driven technique where a conditional moment of a random variable is specified as an unknown function of the data and estimated by means of a kernel or any other weighting scheme on the data.

Random field a scalar random field is defined as a function $m(\omega, x) : \Omega \times A \to R$ such that $m(\omega, x)$ is a random variable for each $x \in A$ where $A \subseteq R^k$.

Sieves the sieves or approximating spaces are approximations to an unknown function, that are dense in the original function space. Sieves can be constructed us-