ESTIMATING THE EQUITY PREMIUM

John Y. Campbell

Estimating the Equity Premium
John Y. Campbell
NBER Working Paper No. 13423
September 2007
JEL No. G12

<u>ABSTRACT</u>

To estimate the equity premium, it is helpful to use finance theory: not the old-fashioned theory that efficient markets imply a constant equity premium, but theory that restricts the time-series behavior of valuation ratios, and that links the cross-section of stock prices to the level of the equity premium. Under plausible conditions, valuation ratios such as the dividend-price ratio should not have trends or explosive behavior. This fact can be used to strengthen the evidence for predictability in stock returns. Steady-state valuation models are also useful predictors of stock returns given the high degree of persistence in valuation ratios and the difficulty of estimating free parameters in regression models for stock returns. A steady-state approach suggests that the world geometric average equity premium was almost 4% at the end of March 2007, implying a world arithmetic average equity premium somewhat above 5%. Both valuation ratios and the cross-section of stock prices imply that the equity premium fell considerably in the late 20th Century, but has risen modestly in the early years of the 21st Century.

John Y. Campbell
Morton L. and Carole S.
Olshan Professor of Economics
Department of Economics
Littauer Center 213
Cambridge, MA 02138
and NBER
john_campbell@harvard.edu

1 Introduction

What return should investors expect the stock market to deliver, above the interest rate on a safe short-term investment? In other words, what is a reasonable estimate of the equity premium?

This question is a basic one for investors who must decide how to allocate their portfolios to safe and risky assets. In the academic world, it has for over three decades played a central role in the development of asset pricing theory and financial econometrics. In the 1960's and 1970's, the efficient market hypothesis was interpreted to mean that the true equity premium was a constant. Investors might update their estimates of the equity premium as more data became available, but eventually these estimates should converge to the truth. This viewpoint was associated with the use of historical average excess stock returns to forecast future returns.

In the early 1980's, a number of researchers reported evidence that excess stock returns could be predicted by regressing them on lagged financial variables. In particular, valuation ratios that divide accounting measures of cash flow by market valuations, such as the dividend-price ratio, earnings-price ratio, or smoothed earnings-price ratio, appeared to predict returns. Value-oriented investors in the tradition of Graham and Dodd (1934) had always asserted that high valuation ratios are an indication of an undervalued stock market and should predict high subsequent returns, but these ideas did not carry much weight in the academic literature until authors such as Rozeff (1984), Fama and French (1988), and Campbell and Shiller (1988a,b) found that valuation ratios are positively correlated with subsequent returns and that the implied predictability of returns is substantial at longer horizons. Around the same time, several papers pointed out that yields on short- and long-term Treasury and corporate bonds are correlated with subsequent stock returns (Fama and Schwert 1977, Keim and Stambaugh 1986, Campbell 1987, Fama and French 1989).

These results suggested that the equity premium is not a constant number that can be estimated ever more precisely, but an unknown state variable whose value must be inferred at each point in time on the basis of observable data. Meanwhile research in asset pricing theory made financial economists more comfortable with the idea that the equity premium can change over time even in an efficient market with rational investors, so that a time-varying equity premium does not necessarily require abandonment of the traditional paradigm of financial economics for a behavioral or inefficient-markets alternative. Campbell and Cochrane (1999), for example, showed

that rational investors with habit formation preferences might become more averse to volatility in consumption and wealth, driving up the equilibrium equity premium, when the economy is weak.

During the 1990's, research continued on regressions predicting stock returns from valuation ratios (Kothari and Shanken 1997, Lamont 1998, Pontiff and Schall 1998) and interest rates (Hodrick 1992). However the 1990's also saw challenges to the new view that valuation ratios predict stock returns.

A first challenge came from financial econometricians, who began to express concern that the apparent predictability of stock returns might be spurious. Many of the predictor variables in the literature are highly persistent: Nelson and Kim (1993) and Stambaugh (1999) pointed out that persistence leads to biased coefficients in predictive regressions if innovations in the predictor variable are correlated with returns (as is strongly the case for valuation ratios, although not for interest rates). Under the same conditions the standard t-test for predictability has incorrect size (Cavanagh, Elliott, and Stock 1995). These problems are exacerbated if researchers are data mining, considering large numbers of variables and reporting only those results that are apparently statistically significant (Foster, Smith, and Whaley 1997, Ferson, Sarkissian, and Simin 2003). An active recent literature discusses alternative econometric methods for correcting the bias and conducting valid inference (Cavanagh, Elliott, and Stock 1995, Lewellen 2004, Torous, Valkanov, and Yan 2004, Campbell and Yogo 2006, Jansson and Moreira 2006, Polk, Thompson, and Vuolteenaho 2006, Ang and Bekaert 2007, Cochrane 2007).

A second challenge was posed by financial history. In the late 1990's valuation ratios were extraordinarily low, so regression forecasts of the equity premium became negative (Campbell and Shiller 1998). Yet stock returns continued to be high until after the turn of the millennium. Data from these years were sufficiently informative to weaken the statistical evidence for stock return predictability. Although low returns in the early 2000's have partially restored this evidence, Goyal and Welch (2003, 2007) and Butler, Grullon, and Weston (2005) have argued that overall, the out-of-sample forecasting power of valuation ratios is often worse than that of a traditional model predicting the equity premium using only the historical average of past stock returns.

The ultimate test of any predictive model is its out-of-sample performance. My personal experience using regression models to forecast stock returns in the late 1990's was humbling, although these models were partially vindicated by the stock market decline of the early 2000's. The lesson I draw from this experience is that one is more

likely to predict stock returns successfully if one uses finance theory to reduce the number of parameters that must be freely estimated from the data, and to restrict estimates of the equity premium to a reasonable range.

In the next section of this paper I show how finance theory can be used if one believes that valuation ratios, in particular the dividend-price ratio, are stationary around a constant mean. Even under stationarity, the persistence of valuation ratios has led researchers to concentrate on situations where valuation ratios have a root that is close to unity. In section 3 I discuss the limiting case where one believes that the dividend-price ratio follows a geometric random walk. I show that this case allows an even larger role for theory: it implies that one should forecast returns by adding a growth estimate to the dividend-price ratio, in the manner of the classic Gordon growth model. I argue that this approach has generated successful out-of-sample forecasts historically, and is likely to do so in the future as well. In section 4 I apply this methodology to estimate the current equity premium for Canada, for the US, and for the world stock market as a whole. In section 5 I briefly discuss how finance theory can be used to predict the equity premium from the cross-section of stock prices. Section 6 concludes.

2 Regression-based return prediction with a stationary dividend-price ratio

When the dividend-price ratio is stationary, a basic tool for analyzing stock returns is the loglinear approximate relation derived by Campbell and Shiller (1988a). This relation says that the log stock return r_{t+1}, the log stock price p_t, and the log dividend d_t approximately satisfy

$$
\begin{aligned}
r_{t+1} &= k + \rho p_{t+1} + (1-\rho)d_{t+1} - p_t \\
&= k + (d_t - p_t) + \Delta d_{t+1} - \rho(d_{t+1} - p_{t+1}),
\end{aligned} \tag{1}
$$

where ρ is a coefficient of loglinearization equal to the reciprocal of one plus the steady-state level of the dividend-price ratio. Thus ρ is slightly smaller than one; for annual US data, $\rho = 0.96$ is a reasonable value given an average dividend-price ratio in the late 20th Century of about 4% or 0.04 in levels. This equation says that proportional changes in stock prices have a larger effect on returns than equal

proportional changes in dividends, because the level of dividends is small relative to the level of prices.

Equation (1) is a difference equation for the log dividend-price ratio. Solving it forward, imposing a condition that there are no explosive bubbles in stock prices, and taking expectations at time t allows us to interpret the dividend-price ratio as

$$d_t - p_t = \frac{k}{1-\rho} + E_t \sum_{j=0}^{\infty} \rho^j [r_{t+1+j} - \Delta d_{t+1+j}]. \tag{2}$$

This formula delivers a number of insights. First, it helps to motivate regressions of stock returns on the log dividend-price ratio. The ratio is a linear combination of discounted expectations of future stock returns and dividend growth. If dividend growth is not too predictable (and there is little direct evidence for long-term dividend predictability in US data), and if the dynamics of discount rates are such that short- and long-term expected stock returns are highly correlated, then the log dividend-price ratio should be a good proxy for the expected stock return over the next period.

Second, equation (2) shows that in the absence of price bubbles, the log dividend-price ratio will be stationary if stock returns and dividend growth are stationary, conditions that seem quite plausible. In particular, if returns and dividend growth rates do not have time trends, then the log dividend-price ratio will not have a time trend either. (This model cannot be used to say what would happen if there were time trends in returns and dividend growth rates, because such trends would invalidate the linear approximation (1).) Third, however, persistent variation in returns and dividend growth rates can lead to persistent variation in the log dividend-price ratio even if that ratio is stationary.

The effect of persistence on predictive regressions has been highlighted by Stambaugh (1999). Stambaugh discusses the two-equation system

$$r_{t+1} = \alpha + \beta x_t + u_{t+1}, \tag{3}$$
$$x_{t+1} = \mu + \phi x_t + \eta_{t+1}, \tag{4}$$

where x_t can be any persistent predictor variable but attention focuses on the level or log of the dividend-price ratio.

OLS estimates of equation (3) in 20th Century US data, with the log dividend-price ratio $x_t = d_t - p_t$ as the explanatory variable and the annualized stock return as the dependent variable, tend to deliver estimates in the range 0.1 to 0.2. An estimate of 0.04, the historical average level of the dividend-price ratio, would imply that around the average, a percentage point increase in the level of the dividend-price ratio increases the expected stock return by one percentage point. The OLS estimates imply a sensitivity of the return to the dividend-price ratio that is several times greater than this. They imply that when the dividend-price ratio is unusually high, it tends to return to normal through increases in prices that magnify the effect on stock returns. Campbell and Shiller (1998) emphasize this pattern in the historical data.

To understand Stambaugh's concern about persistence, define

$$\gamma = \frac{\sigma_{u\eta}}{\sigma_\eta^2}. \tag{5}$$

The coefficient γ is the regression coefficient of return innovations on innovations to the predictor variable. In the case where the explanatory variable is the log dividend-price ratio, γ is negative because rising stock prices tend to be associated with a falling dividend-price ratio. More precisely, dividend growth is only weakly correlated with and much less volatile than stock returns, so from equation (1) γ is about $-\rho$, that is, slightly greater than -1.

Stambaugh points out that the bias in estimating the coefficient β is γ times the bias in estimating the persistence of the predictor variable, ϕ:

$$\mathrm{E}[\widehat{\beta} - \beta] = \gamma \mathrm{E}[\widehat{\phi} - \phi]. \tag{6}$$

This is significant because it has been understood since the work of Kendall (1954) that there is downward bias in estimates of ϕ of about $-(1 + 3\phi)/T$, where T is the sample size, primarily resulting from the fact that x_t has an unknown mean that must be estimated. With a highly persistent predictor variable and γ slightly greater than -1, the Stambaugh bias in $\widehat{\beta}$ is almost $4/T$. With 50 years of data the bias is almost 0.08, substantial relative to the OLS estimates discussed above.

Recent responses to Stambaugh's critique have all used theory in one way or another. Lewellen (2004) first writes an expression for the bias conditional on the estimated persistence $\widehat{\phi}$ and the true persistence ϕ:

$$\mathrm{E}[\widehat{\beta} - \beta \,|\, \widehat{\phi}, \, \phi] = \gamma[\widehat{\phi} - \phi]. \tag{7}$$

At first sight this expression does not seem particularly useful because we do not know the true persistence coefficient. However, Lewellen argues on the basis of theory that ϕ cannot be larger than one—the dividend-price ratio is not explosive—so the largest bias occurs when $\phi = 1$. He proposes the conservative approach of adjusting the estimated coefficient using this worst-case bias:

$$\widehat{\beta}_{adj} = \widehat{\beta} - \gamma(\widehat{\phi} - 1). \tag{8}$$

In the data, the log dividend-price ratio appears highly persistent. That is, $\widehat{\phi}$ is close to one; Lewellen reports a monthly estimate of 0.997 for the period 1946-2000, or about 0.965 on an annual basis. Lewellen's bias adjustment is therefore about 0.035, much smaller than Stambaugh's bias adjustment for a 50-year sample and somewhat smaller whenever the sample size is less than 114 years. Lewellen argues that stock returns are indeed predictable from the log dividend-price ratio, almost as much so as a naive researcher, unaware of Stambaugh's critique, might believe. Another way to express Lewellen's point is that data samples with spurious return predictability are typically samples in which the log dividend-price ratio appears to mean-revert more strongly than it truly does. In the historical data, the log dividend-price ratio has a root very close to unity—it barely seems to mean-revert at all—and thus we should not expect important spurious predictability in the historical data.

Cochrane (2007) responds to Stambaugh by directing attention to the inability of the log-dividend price ratio to forecast dividend growth. At first sight this response does not seem connected to Lewellen's, but in fact it is closely related. The Campbell-Shiller loglinearization (1) implies that r_{t+1}, Δd_{t+1}, $d_{t+1} - p_{t+1}$, and $d_t - p_t$ are deterministically linked. It follows that if we regress r_{t+1}, Δd_{t+1}, and $d_{t+1} - p_{t+1}$ onto $d_t - p_t$, the coefficients β, β_d, and ϕ are related by

$$\beta = 1 - \rho\phi + \beta_d, \tag{9}$$

where ρ is the coefficient of loglinearization from equation (1).

If we have prior knowledge about ϕ, then β and β_d are linked. For example, if $\rho = 0.96$ and we know that $\phi \le 1$, then $\beta_d \le \beta - 0.04$. If $\beta = 0$, then β_d must be negative and less than -0.04. The fact that regression estimates of β_d are close to zero is therefore indirect evidence that $\beta > 0$, in other words that stock returns are predictable—given our prior knowledge, based on theory, that the log dividend-price ratio is not explosive.

Another way to express Cochrane's point is that if the dividend-price ratio fails to predict stock returns, it will be explosive unless it predicts dividend growth. Since the dividend-price ratio cannot be explosive, the absence of predictable dividend growth strengthens the evidence for predictable returns.

Campbell and Yogo (2006) offer a third response to Stambaugh. They point out that if we knew persistence, we could reduce noise by adding the innovation to the predictor variable to the predictive regression, estimating

$$r_{t+1} = \alpha' + \beta x_t + \gamma(x_{t+1} - \phi x_t) + v_{t+1}. \tag{10}$$

The additional regressor, $(x_{t+1} - \phi x_t) = \eta_{t+1}$, is uncorrelated with the original regressor x_t but correlated with the dependent variable r_{t+1}. Thus the regression (10) still delivers a consistent estimate of the original predictive coefficient β, but it does so with increased precision because it controls for some of the noise in unexpected stock returns.

Of course, in practice we do not know the persistence coefficient ϕ, but Campbell and Yogo argue that we can construct a confidence interval for it by inverting a unit root test. By doing this we "de-noise" the return and get a more powerful test. The test delivers particularly strong evidence for predictability if we rule out a persistence coefficient $\phi > 1$ on prior grounds.

A way to understand Campbell and Yogo's results is to recall the challenge posed by the late 1990's. In that period, the dividend-price ratio was low, which led Campbell and Shiller (1998) to predict low stock returns based on a regression like (3). In fact, stock returns remained high until the early 2000's. These high returns were accompanied by falling dividend yields despite the fact that the dividend yield was already below its historical mean. If we believe that the dividend yield was below its true mean and that it should be forecast to return to that mean rather than exploding away from it, then the late 1990's declines in the dividend-price ratio must have been unexpected. Unexpected declines in the dividend-price ratio are associated with unexpected high stock returns, accounting for the poor performance of the basic predictability regression in the late 1990's. The regression (10) corrects for this effect, limiting the negative influence of the late 1990's on the estimated predictive coefficient β.

The econometric issues discussed in this section have little effect on regressions that use nominal interest rates or yield spreads to predict excess stock returns. Although nominal interest rates are highly persistent, their innovations are not strongly

correlated with innovations in stock returns and thus the coefficient γ is close to zero for these variables, implying only a trivial bias in OLS regression estimates. Even papers that are skeptical of stock return predictability from the dividend-price ratio, such as Ang and Bekaert (2007), emphasize the strength of the statistical evidence that interest rates predict stock returns. The challenge in this case is primarily a theoretical one, to understand the economic forces that cause common variation in nominal interest rates and the equity premium.

All the papers discussed above combine prior knowledge with classical statistical methods. It is of course possible to use finance theory in an explicit Bayesian manner. Several recent papers have done this, notably Pastor and Stambaugh (2007) and Wachter and Warusawitharana (2007). Consistent with the results reported here, these papers find that tight priors on the persistence of the predictor variable tend to deliver stronger evidence for predictability of stock returns.

3   Steady-state return prediction

The papers discussed in the previous section address the question of whether the equity premium varies with market valuations, or whether it is constant. Even if one believes that the equity premium is time-varying, however, there remains the important question of how best to estimate it at each point in time. Given the noise in stock returns, equity premium models with multiple free coefficients are hard to estimate and may fail out of sample because of errors in estimating the coefficients. Indeed, Goyal and Welch (2007) argue that almost all the regression models proposed in the recent literature fail to beat the historical sample mean when predicting excess stock returns out of sample.

In response to Goyal and Welch, Campbell and Thompson (2007) propose to use steady-state valuation models to estimate the equity premium. Such models tightly restrict the way in which historical data are used to predict future returns, and Campbell and Thompson find that they work well out of sample. Fama and French (2002) and Pastor, Sinha, and Swaminathan (2007) also use this approach to analyze the equity premium. The approach is analogous to the familiar procedure of forecasting the return on a bond using its yield rather than its historical average return.

The classic steady-state model is the Gordon growth model, named after Canadian economist Myron Gordon. The model describes the level of the dividend-price ratio in a steady state with a constant discount rate and growth rate. Using upper-case letters to denote levels of variables, the Gordon growth model can be written as:

$$\frac{D}{P} = R - G. \tag{11}$$

This formula can be used directly with historical dividend growth rates, but it can also be rewritten in several ways that suggest alternative empirical strategies for forecasting stock returns. First, one can substitute out growth by using the steady-state relation between growth and accounting return on equity,

$$G = (1 - \frac{D}{E})ROE, \tag{12}$$

where D/E is the payout ratio, to obtain a growth-adjusted return forecast

$$\widehat{R}_{DP} = \frac{D}{P} + (1 - \frac{D}{E})ROE. \tag{13}$$

This return forecast is linear in D/P, with a slope coefficient of one and an intercept that is determined by the reinvestment rate and profitability. Importantly, neither the slope coefficient nor the intercept need to be estimated from noisy historical stock returns.

Second, one can restate the model in terms of the earnings-price ratio by using $D/P = (D/E)(E/P)$ to obtain

$$\widehat{R}_{EP} = \left(\frac{D}{E}\right)\frac{E}{P} + (1 - \frac{D}{E})ROE, \tag{14}$$

a payout-ratio-weighted average of the earnings-price ratio and the accounting return on equity. When return on equity equals the expected return, as might be the case in long-run equilibrium, then this implies that $\widehat{R}_{EP} = E/P$.

Finally, one can rewrite the model in terms of the book-market ratio. Since $E/P = (B/M)ROE$,

$$\widehat{R}_{BM} = ROE\left[1 + \frac{D}{E}\left(\frac{B}{M} - 1\right)\right]. \tag{15}$$

To use these formulas in practice, one must decide how to combine historical and contemporaneous data on the right hand side variables. Campbell and Thompson (2007) follow Fama and French (2002) by using historical average data on payouts and profitability, but differ from them by using current rather than historical average data on valuation ratios to obtain a return forecast conditional on the market's current valuation level. This procedure assumes that movements in valuation ratios, relative to historical cash flows, are explained by permanent changes in expected returns so that each percentage point increase in the level of the dividend-price ratio generates a percentage point increase in the return forecast. It is a compromise between the view that valuation ratios are driven by changing forecasts of profitability, in which case the implied movements in returns would be smaller, and the view that valuation ratios are driven by *temporary* changes in discount rates, in which case the implied return movements would be larger as discussed in the previous section.

Campbell and Thompson evaluate the out-of-sample performance of these models and several other variants over the period 1927-2005 and subsamples with breakpoints at 1956 and 1980. They find that steady-state valuation models typically perform better when more theoretical restrictions are imposed, and that they almost always outperform the historical mean return as a predictor of future returns. Dividend-based and earnings-based models, equations (13) and (14), generally appear to be more successful than the book-market model (15). In the next section I illustrate this approach using a model that averages both the dividend-price ratio and the recent history of earnings to generate a return forecast that is a blend of those from (13) and (14).

3.1 The Gordon model with a random walk dividend-price ratio

It may at first sight appear strange that steady-state valuation models based on the Gordon growth model perform well given that they assume constant valuation ratios, while in the data valuation ratios vary in a highly persistent manner. It turns out, however, that a variant of the Gordon growth model can be derived using the assumption that the log dividend-price ratio follows a random walk. Under this assumption the Campbell-Shiller loglinear model, used in the previous section, breaks down because the dividend-price ratio has no fixed mean around which to take a loglinear approximation. However in this case a suitable version of the original Gordon

growth model is available to take the place of the Campbell-Shiller model.

To show this I assume, as in the Gordon growth model, that the dividend is known one period in advance. Then we can write

$$\frac{D_{t+1}}{P_t} = \exp(x_t), \tag{16}$$

where x_t now denotes the log dividend-price ratio using a forward or indicated dividend rather than a historical dividend. I assume that x_t follows a random walk:

$$x_t = x_{t-1} + \varepsilon_t. \tag{17}$$

Since the dividend growth rate is known one period in advance, I can write

$$\frac{D_{t+1}}{D_t} = 1 + G_t = \exp(g_t). \tag{18}$$

Finally, I assume that x_{t+1} and g_{t+1} are conditionally normal given time t information.

The definition of the stock return implies that

$$
\begin{aligned}
1 + R_{t+1} &= \frac{P_{t+1} + D_{t+1}}{P_t} = \frac{D_{t+1}}{P_t} + \frac{D_{t+2}}{D_{t+1}}\frac{D_{t+1}}{P_t}\left(\frac{D_{t+2}}{P_{t+1}}\right)^{-1} \\
&= \exp(x_t)\left[1 + \exp(g_{t+1} - x_{t+1})\right].
\end{aligned} \tag{19}
$$

The conditionally expected stock return can be calculated using the formula for the conditional expectation of lognormally distributed random variables and the martingale property that $E_t x_{t+1} = x_t$:

$$
\begin{aligned}
E_t(1 + R_{t+1}) &= \exp(x_t)\left[1 + E_t \exp(g_{t+1} - x_{t+1})\right] \\
&= \exp(x_t)[1 + \exp(E_t g_{t+1} - x_t + \sigma_g^2/2 + \sigma_x^2/2 - \sigma_{gx})] \\
&= \frac{D_{t+1}}{P_t} + \exp(E_t g_{t+1})\exp(\mathrm{Var}_t(p_{t+1} - p_t)/2).
\end{aligned} \tag{20}
$$

Finally, the right hand side of (20) can be approximated using the facts that for small y, $\exp(y) \approx 1 + y$, and that unexpected log stock returns are approximately

11

equal to unexpected changes in log stock prices:

$$E_t(1 + R_{t+1}) \approx \frac{D_{t+1}}{P_t} + \exp(E_t g_{t+1}) + \frac{1}{2}\text{Var}_t(r_{t+1}). \tag{21}$$

This equation expresses the expected stock return as the level of the dividend yield, plus geometric average dividend growth, plus one-half the variance of stock returns. In the original Gordon model, $\sigma_x^2 = 0$ so the variance of stock returns equals the variance of dividend growth. Since arithmetic average dividend growth equals geometric average dividend growth plus one-half the variance of dividend growth, in this case we get the original Gordon formula that the arithmetic average stock return equals dividend yield plus arithmetic average dividend growth.

If one subtracts half the variance of stock returns from each side of (20), one finds that the geometric average stock return equals the level of the dividend-price ratio plus the geometric average of dividend growth. Under the assumptions of the original Gordon model, the geometric implementation of the model is equivalent to an arithmetic implementation because stock returns and dividend growth have the same variance so their geometric and arithmetic averages differ by the same amount. In the data, however, returns are much more volatile so the geometric implementation and the arithmetic implementation are different. The analysis here shows that the geometric implementation is correct. Interestingly, this is exactly the way in which the model is used by Siegel (1994).

4 What is the equity premium today?

I now use a version of the above methodology, starting from equation (14), to estimate the equity premium. Following the previous discussion, I first estimate the conditional geometric average stock return, then subtract the real interest rate to get an equity premium number, and finally discuss the adjustment that is needed to convert from a geometric average to an arithmetic average equity premium. I look at data for the world as a whole (measured using the Morgan Stanley Capital International all-world index), and also for the US and Canada, over the period from 1982 through the end of March 2007.

Figure 1 shows that for all three indices smoothed earnings-price ratios, with earnings smoothed over three years to eliminate cyclical noise, have fallen dramatically

since the early 1980's and have been in the 3% to 5% range for the last ten years. During the same period, however, Figure 2 shows that profitability has increased from a long-run historical average of around 6% to much higher values around 10%. Meanwhile payout ratios have fluctuated widely around an average of about 50%.

In constructing a return forecast, it is desirable to combine historical earnings with some forward-looking measure of earnings. One possibility is to use analysts' earnings forecasts (Pastor, Sinha, and Swaminathan 2007), another is to use dividends. I average historical earnings, smoothed over three years, and the current dividend, divided by the payout rate, to construct a forward-looking measure of permanent earnings that can be used in equation (14).

Putting these numbers together, an earnings-based estimate of the real return on US equities, assuming constant 6% real profitability and a 50% payout rate, was about 9% in the early 1980's and fell to just above 4% in the year 2000. Since then it has increased to slightly over 5%. This estimate assumes that profitability and payouts are best forecast to be constant; alternatively, if one uses the 3-year moving average of profitability illustrated in Figure 2, and a similar 3-year moving average of the payout ratio, the current real return estimate increases by almost 4% to 9%, reflecting the high recent profitability and low payout ratios of US corporations. At the world level, the current real return number is comparable to the US number if a fixed profitability estimate is used, but the adjustment for recent profitability and payouts is much smaller, only slightly above 2%. The Canadian real return number is also very similar to that in the US on the basis of fixed profitability, but lower Canadian profitability and higher payouts in the last few years imply that the use of recent data increases the estimated real return by less than 2%.

To convert these numbers into estimates of the equity premium, one needs to subtract a safe real interest rate. Figure 3 plots real yields on inflation-indexed bonds in three large markets, the UK, the US, and Canada. The figure shows that the average real yield on inflation-indexed bonds across the three countries was about 3.5% in the 1990's but fell below 2% in the early 2000's. By the end of March 2007, it had recovered to just over 2%.

The implied current equity premium, assuming constant profitability and payouts, is just over 3%: 3.3% for the world as a whole, 3.2% for the US, and 3.1% for Canada. If instead one uses recent profitability and payouts, the current equity premium is 5.7% for the world as a whole, a startling 6.9% for the US, and 5.0% for Canada. Figures 4, 5, and 6 illustrate the history of the equity premium in the world, the US,

and Canada under these two alternative assumptions.

Obviously a key question is whether the high profitability of global, and particularly US, corporations can be expected to continue. On the one hand, globalization has increased the supply of labor relative to capital, reducing wage pressure and increasing profitability; on the other hand, profitability has been increased by favorable business cycle and political conditions that may not persist. Historically, profitability has shown temporary fluctuations and low payout rates (high reinvestment rates) have predicted declining profitability. Also, equity premium estimates based on current profitability and payout rates have been highly volatile, even turning negative on occasion. For both these reasons it seems wise to place considerably more weight on long-term averages than on recent data. If one puts a weight of 0.75 on the long-term average, with 0.25 on the recent data, the implied equity premium at the end of March 2007 is in the range 3.6% to 4.1%: 3.9% in the world as a whole, 4.1% in the US, and 3.6% in Canada. This number is a geometric average equity premium; for an arithmetic average, one should add one-half the variance of stock returns, or almost 1.3% if stock returns have a conditional standard deviation of 16%. The resulting arithmetic equity premium numbers are in the range 4.9% to 5.4%. Note that the equity premium is this high in large part because the safe real interest rate has declined over the past decade as illustrated in Figure 3.

These numbers are lower than historical average excess stock returns reported by Dimson, Marsh, and Staunton (2006). Using data for the period 1900-2005, Dimson, Marsh, and Staunton report geometric average equity premia of 4.7% for the world as a whole, 5.5% for the US, and 4.5% for Canada. The difference reflects two facts. First, historical average returns have been driven up by declining valuation ratios; this effect cannot be expected to continue in the future because valuation ratios should not have trends, a point emphasized by Fama and French (2002). Second, historical average returns were obtained by investors who paid lower stock prices and thus benefited from higher dividend-price ratios.

It is interesting to note that chief financial officers of major corporations, surveyed by Graham and Harvey (2007), have modest expectations of the equity premium that imply they do not expect recent profitability to continue. Their median estimate of the geometric average US equity premium at the end of November 2006 was 3.4%, much closer to the constant-profitability number reported here than to the recent-profitability number, and far below the historical average equity premium.

14

5 Return prediction with cross-sectional variables

Finance theory can also be used to predict excess stock returns using information in the cross-section of stock prices. This is valuable both to corroborate the predictions from aggregate valuation ratios, and possibly as a way to pick up higher-frequency components of the equity premium that may be missed by a steady-state approach.

Polk, Thompson, and Vuolteenaho (2006) argue that if the Capital Asset Pricing Model (CAPM) is true, then a high equity premium implies low prices for stocks that have high betas with the aggregate market index. That is, high-beta stocks should be value stocks with low ratios of market prices to accounting measures of fundamental value. Reversing the argument, value stocks should tend to have high betas. This was true in the mid-20th Century, roughly from the 1930's through the 1950's, but in recent decades growth stocks have had higher betas than value stocks (Franzoni 2006). Polk, Thompson, and Vuolteenaho argue that this change in cross-sectional stock pricing reflects a decline in the equity premium. They construct a predictor of the aggregate market return, based on the relative pricing of high- and low-beta stocks, and show that it correlates well with the smoothed earnings-price ratio except in the early 1980's when inflation may have distorted the relationship.

It is possible to push this idea even further, exploiting the fact that the CAPM may not fully describe the cross-section of stock returns when returns are predictable in the time series. Merton (1973) developed an intertemporal CAPM (ICAPM) that showed that in the presence of time-varying expected returns, long-lived investors care not only about shocks to their wealth but also about shocks to the expected return on wealth. Intuitively, they value wealth not for its own sake but for the consumption stream it can provide; thus they want to hedge against declines in the rate of return just as much as against declines in market value. Campbell (1993) implemented this idea using a vector autoregression (VAR) to break market movements into permanent movements driven by news about cash flows and temporary movements driven by news about discount rates. Long-lived investors are more concerned about the former than about the latter. Thus stocks that covary with cash-flow news should have higher average returns than stocks that covary with discount-rate news, controlling for betas with the overall market return.

One of the main deviations from the CAPM in recent decades has been the value effect, the high average returns that value stocks have delivered despite their low market betas. If the ICAPM is to explain the value effect, it must be that value stocks

covary with cash-flow news while growth stocks covary with discount-rate news. This implies that a moving average of past excess returns on growth stocks should be a good predictor of aggregate stock returns.

The value spread, the relative valuation of value and growth stocks (normally measured as the difference between the log book-market ratios of these two types of stocks) is one possible summary of past excess returns on growth stocks. Eleswarapu and Reinganum (2004) find that the value spread for small stocks predicts the aggregate market return, and Campbell and Vuolteenaho (2004) use the same variable in a VAR model to estimate and test the ICAPM. They find that the ICAPM explains the average returns of value and growth stocks much better than does the standard CAPM. Cohen, Polk, and Vuolteenaho (2006) and Campbell, Polk, and Vuolteenaho (2007) explore the robustness of these results, using both VAR-based and direct measures of cash-flow and discount-rate news. Empirically, the effect of including the small-stock value spread in a model of the equity premium is to lower the estimated equity premium at the turn of the millennium, when growth stocks were abnormally expensive relative to value stocks, and to increase it in 2006 and early 2007, when growth stocks were abnormally cheap.

All this work relies on theoretically motivated, but not fully restricted, time-series models of the aggregate market return. A natural next step is to use the theoretical restrictions of the ICAPM to jointly estimate a time-series model of the aggregate market return and a cross-sectional model of average stock returns. Campbell (1996) was an early implementation of this approach, but that paper did not find systematic deviations from the CAPM because it did not use the information in the relative prices of growth and value stocks. Recent research suggests that with the proper information variables and test assets, cross-sectional information can play an important role in a jointly estimated model of the equity premium.

6   Conclusion

In this paper I have tried to illustrate the usefulness of finance theory for statistical analysis of stock returns, and in particular for estimation of the equity premium. The literature on this topic is vast, and inevitably I have neglected some important aspects. Five omissions deserve special mention.

First, I have not reviewed the simple but important point that excess stock returns should be difficult to predict, because highly predictable excess returns would imply extremely large profits for market-timing investors. Campbell and Thompson (2007) explore the mapping from R^2 statistics in predictive regressions to profits and welfare gains for market timers. The basic lesson is that investors should be suspicious of predictive regressions with high R^2 statistics, asking the old question "If you're so smart, why aren't you rich?"

Second, I have confined attention to short-term predictive regressions and have not considered direct forecasts of long-horizon returns. It has been known since Fama and French (1988) that long-horizon regressions often have higher R^2 statistics than short-horizon regressions, but their statistical properties are controversial. Campbell (2001) and Cochrane (2007) argue that in certain circumstances, long-horizon regressions can have superior power to detect predictability when in fact it exists.

Third, I have not discussed recent work that uses finance theory to infer the equity premium from the actions of market participants. Lettau and Ludvigson (2001), for example, argue that the level of consumption in relation to aggregate financial wealth and labor income reveals consumers' expectations of future stock returns. In a similar spirit Baker and Wurgler (2000) use the financing decisions of corporations to infer corporate managers' beliefs about expected stock returns.

Fourth, I have presented estimates of the equity premium without discussing the uncertainty of these estimates. I have suggested that finance theory can reduce our uncertainty about the equity premium, but a more formal Bayesian analysis would be needed to quantify this effect.

Finally, I have not attempted to review the important body of empirical work on the estimation of stock market risk. Mechanically, the volatility of stock returns determines the wedge between geometric and arithmetic average stock returns. Economically, both risk and return matter to investors, and it is plausible that changing risk is one factor that drives the changing equity premium. Merton (1980), Campbell (1987), French, Schwert, and Stambaugh (1987), Harvey (1989), and Glosten, Jagannathan, and Runkle (1993) are a few of the earlier papers that explore this relation. Recent contributions by Ghysels, Santa-Clara, and Valkanov (2005) and Pastor, Sinha, and Swaminathan (2007) find that the equity premium does covary positively with estimated risk, but that this effect does not explain the predictability of stock returns from valuation ratios or interest rates.

Despite the size and complexity of the literature on the equity premium, it has a simple unifying theme. Campbell, Lo, and MacKinlay (1997) argue that "what distinguishes financial economics is the central role that uncertainty plays in both financial theory and its empirical implementation". Theory tells us why stock returns are so hard to predict. But it also holds out the promise of better prediction than we can hope to achieve by purely statistical forecasting methods.

References

Ang, Andrew, and Geert Bekaert (2007) 'Stock return predictability: Is it there?', *Review of Financial Studies* 20, 651–707

Baker, Malcolm, and Jeffrey Wurgler (2000) 'The equity share in new issues and aggregate stock returns', *Journal of Finance* 55, 2219–2257

Boudoukh, Jacob, Roni Michaely, Matthew Richardson, and Michael Roberts (2007) 'On the importance of measuring payout yield: Implications for empirical asset pricing', *Journal of Finance* 62, 877–916

Butler, Alexander W., Gustavo Grullon, and James P. Weston (2005) 'Can managers forecast aggregate market returns?', *Journal of Finance* 60, 963–986

Campbell, John Y. (1987) 'Stock returns and the term structure', *Journal of Financial Economics* 18, 373–399

— (1993) 'Intertemporal asset pricing without consumption data', *American Economic Review* 83, 487–512

— (1996) 'Understanding risk and return', *Journal of Political Economy* 104, 298–345

— (2001) 'Why long horizons? A study of power against persistent alternatives', *Journal of Empirical Finance* 8, 459–491

Campbell, John Y., and John H. Cochrane (1999) 'By force of habit: A consumption-based explanation of aggregate stock market behavior', *Journal of Political Economy* 107, 205–251

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay (1997) *The Econometrics of Financial Markets* (Princeton, NJ: Princeton University Press)

Campbell, John Y., Christopher Polk, and Tuomo O. Vuolteenaho (2007) 'Growth or glamour? Fundamentals and systematic risk in stock returns', working paper, Harvard University and London School of Economics

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

— (1988b), 'Stock prices, earnings, and expected dividends', *Journal of Finance* 43, 661–676

— (1998), 'Valuation ratios and the long-run stock market outlook', *Journal of Portfolio Management* 24(2), 11–26

Campbell, John Y., and Samuel B. Thompson (2007) 'Predicting excess stock returns out of sample: Can anything beat the historical average?', forthcoming *Review of Financial Studies*

Campbell, John Y., and Tuomo O. Vuolteenaho (2004) 'Bad beta, good beta', *American Economic Review* 94, 1249–1275

Campbell, John Y., and Motohiro Yogo (2006) 'Efficient tests of stock return predictability', *Journal of Financial Economics* 81, 27–60

Cavanagh, Christopher L., Graham Elliott, and James H. Stock (1995) 'Inference in models with nearly integrated regressors', *Econometric Theory* 11, 1131 –1147

Cochrane, John H. (2007) 'The dog that did not bark: A defense of return predictability', forthcoming *Review of Financial Studies*

Cohen, Randolph B., Christopher Polk, and Tuomo O. Vuolteenaho (2006) 'The price is (almost) right', working paper, Harvard University and Northwestern University

Dimson, Elroy, Paul Marsh, and Michael Staunton (2006) 'The worldwide equity premium: a smaller puzzle', working paper, London Business School

Eleswarapu, Venkat R. and Marc R. Reinganum (2004) 'The predictability of aggregate stock market returns: Evidence based on glamour stocks', *Journal of Business* 77, 275–294

Fama, Eugene F., and Kenneth R. French (1988) 'Dividend yields and expected stock returns', *Journal of Financial Economics* 22, 3–25

— (1989) 'Business conditions and expected returns on stocks and bonds', *Journal of Financial Economics* 25, 23–49

— (2002) 'The equity premium', *Journal of Finance* 57, 637–659

Fama, Eugene F., and G. William Schwert (1977) 'Asset returns and inflation', *Journal of Financial Economics* 5, 115–146

Ferson, Wayne E., Sergei Sarkissian, and Timothy T. Simin (2003) 'Spurious regressions in financial economics?', *Journal of Finance* 58, 1393–1413.

Foster, F. Douglas, Tom Smith, and Robert E. Whaley (1997) 'Assessing goodness-of-fit of asset pricing models: the distribution of the maximal R^2', *Journal of Finance* 52, 591–607

Franzoni, Francesco (2006) 'Where is beta going? The riskiness of value and small stocks', working paper, HEC School of Management

French, Kenneth R., G. William Schwert, and Robert F. Stambaugh (1987) 'Expected stock returns and volatility', *Journal of Financial Economics* 19, 3–29

Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov (2005) 'There is a risk-return tradeoff after all', *Journal of Financial Economics* 76, 509–548

Glosten, Lawrence R., Ravi Jagannathan, and David E. Runkle (1993) 'On the relation between the expected value and the volatility of the nominal excess return on stocks', *Journal of Finance* 48, 1779–1801

Goyal, Amit, and Ivo Welch (2003) 'Predicting the equity premium with dividend ratios', *Management Science* 49, 639–654

— (2007) 'A comprehensive look at the empirical performance of equity premium prediction', forthcoming *Review of Financial Studies.*

Graham, Benjamin, and David L. Dodd (1934) *Security Analysis*, 1st ed. (New York, NY: McGraw-Hill)

Graham, John, and Campbell Harvey (2007) 'The equity risk premium in January 2007: Evidence from the global CFO outlook survey', working paper, Duke University.

Harvey, Campbell (1989) 'Time-varying conditional covariances in tests of asset pricing models', *Journal of Financial Economics* 24, 289–317

Hodrick, Robert J. (1992) 'Dividend yields and expected stock returns: Alternative procedures for inference and measurement', *Review of Financial Studies* 5, 257–286

Jansson, Michael, and Marcelo J. Moreira (2006) 'Optimal inference in regression models with nearly integrated regressors', *Econometrica* 74, 681–715

Keim, Donald B., and Robert F. Stambaugh (1986) 'Predicting returns in the stock and bond markets', *Journal of Financial Economics* 17, 357–390

Kendall, Maurice G. (1954) 'Note on bias in the estimation of autocorrelation', *Biometrika* 41, 403–404

Kothari, S.P., and Jay Shanken (1997) 'Book-to-market, dividend yield, and expected market returns: A time-series analysis', *Journal of Financial Economics* 44, 169–203

Lamont, Owen (1998) 'Earnings and expected returns', *Journal of Finance* 53, 1563–1587

Lettau, Martin, and Sydney Ludvigson (2001) 'Consumption, aggregate wealth, and expected stock returns', *Journal of Finance* 56, 815–849

Lewellen, Jonathan (2004) 'Predicting returns with financial ratios', *Journal of Financial Economics* 74, 209–235

Merton, Robert C. (1973) 'An intertemporal capital asset pricing model', *Econometrica* 41, 867–887

— (1980) 'On estimating the expected return on the market: an exploratory investigation', *Journal of Financial Economics* 8, 323–361

Nelson, Charles, and M. Kim (1993) 'Predictable stock returns: The role of small sample bias', *Journal of Finance* 48, 641–661

Pastor, Lubos, Meenakshi Sinha, and Bhaskaran Swaminathan (2007) 'Estimating the intertemporal risk-return tradeoff using the implied cost of capital', forthcoming *Journal of Finance*

Pastor, Lubos, and Robert F. Stambaugh (2007) 'Predictive systems: living with imperfect predictors', NBER working paper 12814.

Polk, Christopher, Samuel Thompson, and Tuomo Vuolteenaho (2006) 'Cross-sectional forecasts of the equity premium', *Journal of Financial Economics* 81, 101–141

Pontiff, Jeffrey, and Lawrence D. Schall (1998) 'Book-to-market ratios as predictors of market returns', *Journal of Financial Economics* 49, 141–160

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

Siegel, Jeremy (1994) *Stocks for the Long Run* (New York, NY: McGraw-Hill)

Stambaugh, Robert F. (1999) 'Predictive regressions', *Journal of Financial Economics* 54, 375–421

Torous, Walter, Rossen Valkanov, and Shu Yan (2004) 'On predicting stock returns with nearly integrated explanatory variables', *Journal of Business* 77, 937–966

Wachter, Jessica A., and Missaka Warusawitharana (2007) 'Predictable returns and asset allocation: Should a skeptical investor time the market?', NBER working paper 13165
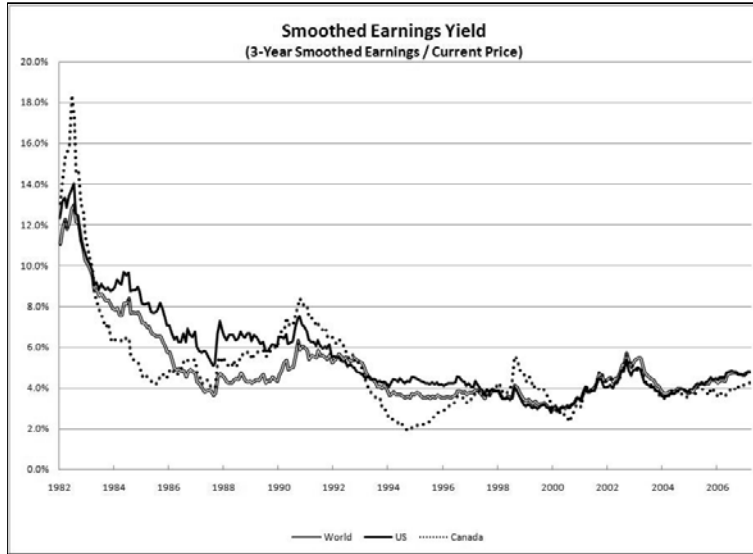
Figure 1: Three-year smoothed earnings-price ratios in the world, the US, and Canada
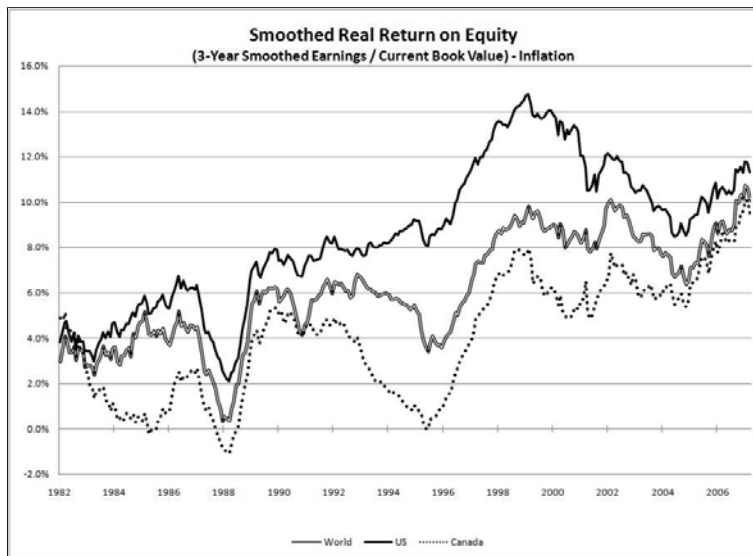

Figure 2: Three-year smoothed profitability in the world, the US, and Canada
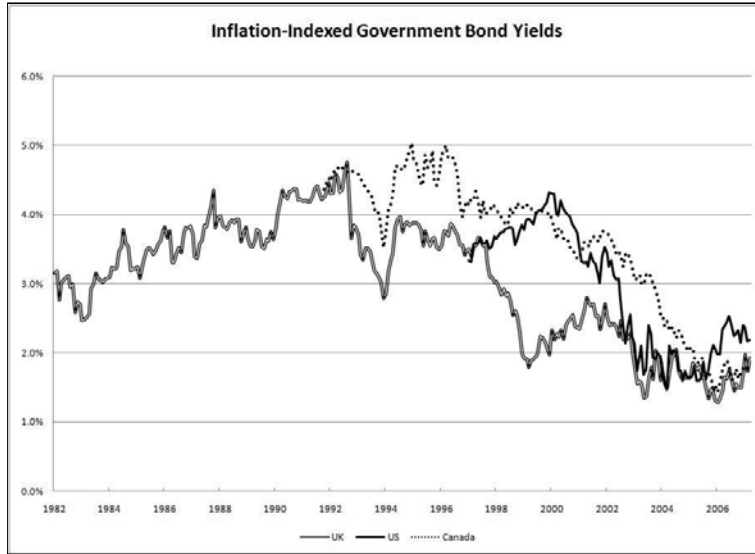
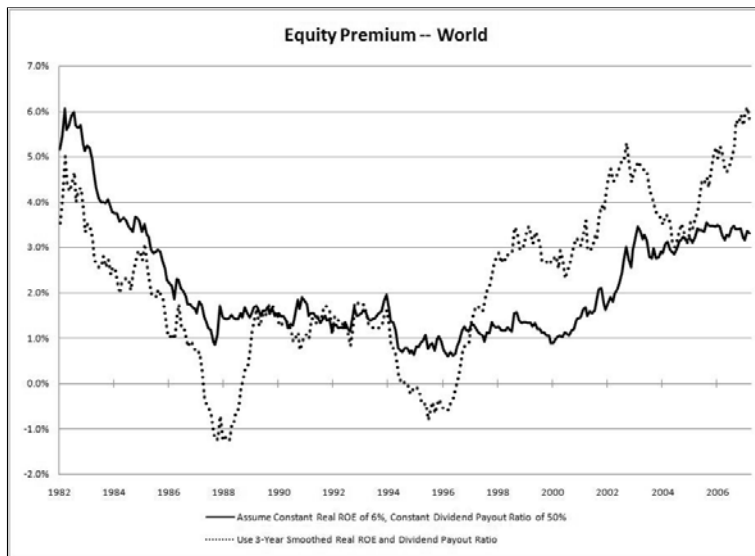Figure 3: Long-term real interest rates in the UK, the US, and Canada
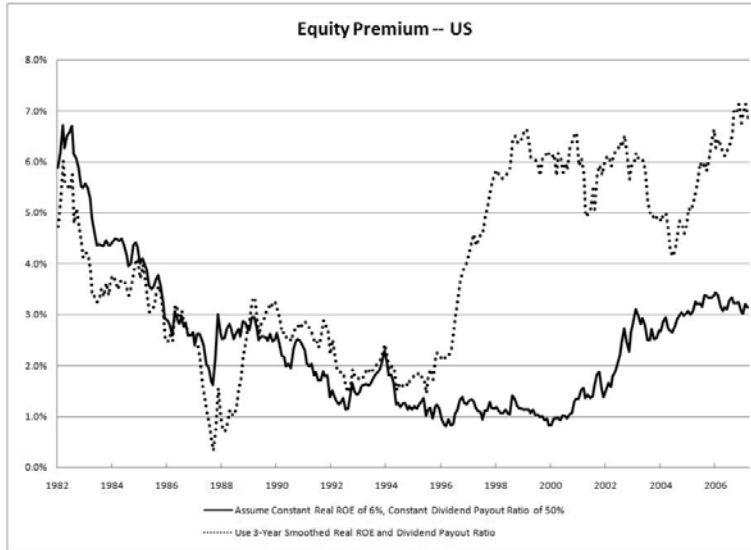

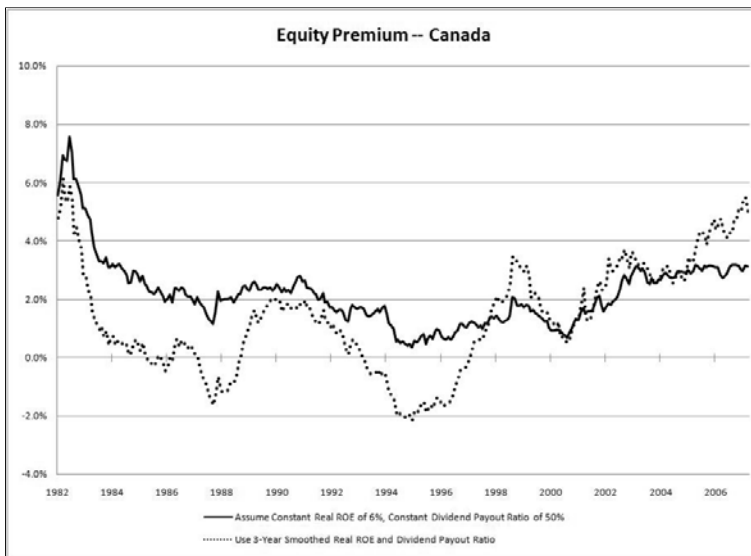Figure 4: The world equity premium since 1982

Figure 5: The US equity premium since 1982

Figure 6: The Canadian equity premium since 1982