Estimating the *Ex Ante* Equity Premium

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

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

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

November 2006

Estimating the *Ex Ante* Equity Premium

Abstract

Over the past century, US stocks have returned roughly 6% more than risk-free debt, which is an equity premium much higher than implied by standard economic theory. Estimates from several studies suggest the equity premium has fallen in recent years. However, it is well known that all equity premium estimates reported to date are imprecisely measured. Many estimates of the *ex ante* equity premium in the literature have confidence bounds as wide as 1000 basis points. Some recent estimates have confidence bounds as narrow as 320 basis points. We aim to narrow further the range of plausible equity premium estimates, taking inspiration from some recent studies that use fundamental information like dividends to estimate the equity premium. Our contribution is based on a new method that simulates the distribution from which interest rates, dividend growth rates, and equity premia are drawn and determines the prices and returns consistent with these distributions. By comparing statistics that arise from our simulations with key financial characteristics of the US economy, including the mean dividend yield, return volatility, and mean return, we are able to narrow substantially the range of plausible equity premium estimates. We consider a broad collection of data generating processes that incorporate features such as a downward trending equity premium, a structural break in the equity premium time-series process, an upward shift in dividend growth rates, sampling uncertainty in generating model parameters, and cross-correlation between interest rates, dividend growth rates, and equity premia. The result of our analysis is by far the most precise equity premium estimate that has been reported in the literature to date. We find that the true *ex ante* equity premium lies within 50 basis points of 3.5%.

Estimating the *Ex Ante* Equity Premium

Financial economic theory is often concerned with the premium that investors demand *ex ante*, when they first decide whether to purchase risky stocks instead of risk-free debt. In contrast, empirical tests of the equity premium often focus on the return investors received *ex post*.[1] It is well known that estimates of the *ex ante* equity premium based on *ex post* data can be very imprecise; such estimates have very wide margins of error, as wide as 1000 basis points in typical studies and 320 basis points in some recent studies. This makes it challenging to employ the estimates for common practical purposes, including evaluating the equity premium puzzle, performing valuation, and conducting capital budgeting. Our goal, therefore, is to develop a more precise estimate of the *ex ante* equity premium. We accomplish this goal by employing powerful simulation techniques that identify a range of values of the *ex ante* equity premium that are consistent with values of several key financial statistics that are actually observed in the US economy, including dividend growth rates, interest rates, Sharpe ratios, price-dividend ratios, volatilities, and of course the *ex post* equity premium. Our results suggest that the true *ex ante* equity premium lies within 50 basis points of 3.5%.

We draw on two relatively new techniques in order to provide a more precise estimate of the equity premium than is currently available. The first technique is the dividend discounting method of fundamental valuation of Donaldson and Kamstra (1996). This technique permits the simulation of fundamental prices, returns, and return volatility for a given *ex ante* equity premium. Donaldson and Kamstra find that if we allow dividend growth rates and discount rates to be time-varying and dependent as well as cross-correlated, the fundamental prices and returns that come out of dividend discounting match observed prices and returns, even during extreme events like the stock market crash of 1929. The second technique is closely related to simulated generalized method of moments (SGMM).[2] SGMM estimates model parameters by using a given model with a given set of parameter values to simulate moments of the data (for instance means or volatilities), measuring

---

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

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

the distance between the simulated moments and the actual data moments, and repeating with new parameter values until the parameter values that minimize the distance are found. An attractive feature of SGMM is that it provides for tests that can reject a model, even at the "best" parameter values, the ones which minimize the distance between the simulated moments and the actual data moments. We exploit the dividend discounting method of Donaldson and Kamstra to generate simulated fundamental prices, dividends, returns, and derivative moments like the mean *ex post* equity premium, mean dividend yield, and return volatility for a given *ex ante* equity premium. We minimize (by choice of the *ex ante* equity premium) the distance between the simulated moments the model produces and the moments observed in US stock markets over the past half century. That is, given various characteristics of the US economic experience (such as low interest rates and high *ex post* equity premia, high Sharpe ratios and low dividend yields, *etc.*), we determine what values of the *ex ante* equity premium are most likely to have generated the observed collection of sample moments.

We consider a broad collection of models, including models with and without trends in the equity premium, with and without breaks in the equity premium, with and without breaks in the dividend growth rate, as well as various autoregressive specifications for dividend growth rates, interest rates, and the equity premium. Virtually every model we consider achieves a minimum distance between the simulated moments and the actual data moments by setting the *ex ante* equity premium between 3% and 4%, typically very close to 3.5%. That is, the equity premium estimate is very close to 3.5% across our models. Further, the range of *ex ante* equity premia that can be supported by the US data for a given model is typically under 100 basis points. Our favored models, those which are based on data generating processes that most closely approximate actual US data, imply the true *ex ante* equity premium is 3.5% plus or minus 50 basis points.

Empirical work investigating the equity premium has consisted of a series of innovations around a common theme: producing a better estimate of the *ex ante* equity premium. Studies typically estimate the equity premium with great imprecision: often a 5% or 6% *ex post* estimate can not be statistically distinguished from an *ex ante* value as low as 1% or as high as 10%. Recent work in the area has included insights such as exploiting dividend yields or earnings yields to provide new, more precise estimates of the return to holding stocks (see Fama and French, 2002, and Jagannathan,

2

McGrattan, and Scherbina, 2000), looking across many countries to account for survivorship issues (see Jorion and Goetzmann, 1999), looking across many countries to decompose the equity premium into dividend growth, price-dividend ratio, dividend yield, and real exchange rate components (see Dimson, Marsh, and Staunton, 2006), modeling equity premium structural breaks and trends in a Bayesian econometric framework (see Pástor and Stambaugh, 2001), or computing out-of-sample forecasts of the distribution of excess returns, allowing for structural breaks which are identified in real time (see Maheu and McCurdy, 2006). Most of this work estimates the *ex ante* equity premium by considering one moment of the data at a time, typically the mean difference between an estimate of the return to holding equity and a risk free rate, though Maheu and McCurdy (2006) consider higher-order moments of the excess return distribution and Pástor and Stambaugh (2001) incorporate return volatility and direction of price movements through their use of priors.

It is the issue of increasing the precision of estimates of the equity premium which motivates this paper. Pástor and Stambaugh (2001), exploiting extra information from return volatility and prices, narrow a two standard deviation confidence interval around the value of the *ex ante* equity premium to plus or minus roughly 280 basis points around a premium estimate of roughly 4.8%, a range that spans about 2% to 7.6%. Fama and French (2002), based on data from 1951 to 2000, provide point estimates of the mean *ex post* equity premium of 4.32% (based on earnings growth rate fundamentals) plus or minus roughly 400 basis points (again, two standard deviations) and of 2.55% (based on dividend growth rate fundamentals) plus or minus roughly 160 basis points: a range of approximately 0.95% to 4.15%. That is, the plausible range of equity premia that emerge from Fama and French's study occupy a confidence bound of 320 to 800 basis points. Claus and Thomas (2001), like Fama and French (2002), make use of fundamental information to form lower estimates of the *ex post* equity premium, but their study covers a shorter time period relative to the Fama and French study – 14 years versus 50 years – yielding point estimates that are subject to at least as much variability as the Fama and French estimates.

The gain in estimation precision we achieve in this paper comes from the intuition behind the gains achieved by Pástor and Stambaugh (2001), Claus and Thomas (2001), Jagannathan, McGrattan, and Scherbina (2000), and Fama and French (2002), that is, exploiting information beyond the simple mean difference between realized equity returns and a risk free rate. We incorporate infor-

3

mation on the US economy's realized dividend yield, return volatility, mean return, *ex post* equity premium, and conditional moments including the correlation of returns with lagged returns (*i.e.* autocorrelation) and return volatility with lagged volatility (*i.e.*, ARCH), using the joint distribution of these statistics and subsets of these statistics to provide formal tests on various values of the *ex ante* equity premium. The moments of the data which we find to be most informative about the *ex ante* equity premium are the mean dividend yield, return volatility, and mean return. While we restrict our attention to sharpening the precision of equity premium estimates for a stock market index, the technique we employ is more broadly applicable to estimating the equity premium of an individual firm.

The remainder of our paper proceeds as follows. The basic methodology used to conduct our simulation approach to estimating equity premia is presented in Section 1, along with important details on estimating the equity premium. (Appendices to the paper provide detailed explanations of the technical aspects of our simulations, including calibration of key model parameters.) In Section 2 we compare univariate financial statistics that arise in our simulations with US market data, including dividend yields, Sharpe ratios, and interest rates. Our results confirm that the simulations generate data broadly consistent with the US market data and, taken one-at-a-time, these financial statistics imply that the *ex ante* equity premium lies in a range much narrower than between 2% and 8%. We determine how much narrower in Section 3 by exploiting the full power of our simulation methodology. We compare joint multivariate distributions of our simulated data with observed US data, yielding a very precise estimate of the equity premium. That is, the range of *ex ante* equity premia that are consistent with the US historical experience is very narrow. Our consideration of a broad collection of possible data generating processes and models lends confidence to the findings. Section 4 concludes.

1   Methodology

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

4