Article:
Freeman, Mark C and Groom, Ben (2015) Using equity premium survey data to estimate future wealth. Review of Quantitative Finance and Accounting. pp. 665-693. ISSN: 0924-865X

https://doi.org/10.1007/s11156-014-0451-7

Noname manuscript No.
(will be inserted by the editor)

Using equity premium survey data to estimate future wealth

Mark C. Freeman · Ben Groom

the date of receipt and acceptance should be inserted later

Abstract We present the first systematic methods for combining different experts' responses to equity premium surveys. These techniques are based on the observation that the survey data are approximately gamma distributed. This distribution has convenient analytical properties that enable us to address three important problems that investment managers must face. First, we construct probability density functions for the future values of equity index tracker funds. Second, we calculate unbiased and minimum least square error estimators of the future value of these funds. Third, we derive optimal asset allocation weights between equities and the risk-free asset for risk-averse investors. Our analysis allows for both herding and biasedness in expert responses. We show that, unless investors are highly uncertain about expert biases or forecasts are very highly correlated, many investment decisions can be based solely on the mean of the survey data minus any expected bias. We also make recommendations for the design of future equity premium surveys.

Keywords Financial surveys · Equity premium · Asset allocation · Gamma distribution
JEL classification G11; G17.

1 Introduction

In recent years a number of surveys have been conducted to solicit informed individuals' views on the forward looking equity premium. For example Welch (2000, with updates in 2001, 2007 and 2009) and Fernandez and del Campo (2010a) ask economics and finance professors; Graham and Harvey (2010) is their latest in a series of studies that has taken data from the Global CFO Outlook Survey; whilst Fernandez and del Campo (2010b) take forecasts from analysts and companies. This stream of research has generated considerable interest amongst both academics and practitioners. In this paper we present the first systematic techniques for optimally combining the conflicting forecasts of the forward looking equity premium presented in any one of these surveys.

Our methods exploit the observation that estimates of the equity premium are, to good approximation, gamma distributed. This draws in parallels with Weitzman's (2001) highly influential use of survey data on the appropriate long term social discount rate for evaluating investments with intergenerational consequences. This distribution has the benefit of a well defined moment generating function (mgf) and established numerical techniques for determining

Address(es) of author(s) should be given

the probability density function (pdf) of the population mean conditional on a finite sample of data.

We use these two features of the gamma distribution to answer questions about the future value of an equity index tracker fund and optimal asset allocation. The fact that a closed-form mgf of the distribution exists allows us to determine unbiased and minimum least square error (MLSE) estimates of the future value, and quantiles of the future value, of the fund. This extends the analysis of Jacquier et al. (2003, 2005) and Kan and Zhou (2009) to forward-looking estimates of the equity premium. That we are able to derive a pdf of the true equity premium based on the survey data also enables us to construct a pdf for the future value of the index tracker fund. These are issues that are of pressing importance for portfolio managers facing asset allocation decisions, actuaries who use stochastic asset-liability models when estimating pension fund solvency, and regulators who prescribe rates of return for illustrating future portfolio values.

We conclude that, unless there is extremely strong herding amongst experts, then any difference between sample and population mean estimates of the equity premium is unlikely to be of any significance for asset allocation and future value problems. This result holds because the standard deviation of survey data is very low; around 1.77% in the case of Welch's 2009 update. With over one hundred responses the standard error associated with the population mean of expert responses is extremely small unless there is very high inter-dependence between forecasts. If, in addition, the decision maker can estimate with reasonable precision the relationship between the population mean of expert responses and the true value of the ex-ante equity premium, then even the presence of uncertain biases does not materially alter the conclusion that, for most cases considered, practical decisions can be made on the basis of the mean of the survey data minus the expected bias alone.

We also make three recommendations for the future design of equity premium surveys. First, researchers should place greater emphasis on obtaining forecasts from a sample of experts who are by their nature heterogeneous rather than increasing the sample size within one particular subsection of the population. Second, including a textual question that asks respondents how they arrived at their answer will help the researcher understand the nature of herding of forecasts. Third, by also asking a question about short-term growth in a variable with low ex-post noise, such as GDP, the extent of analysts' over-optimism (-pessimism) may be revealed, allowing us to understand something of the nature of the bias in the forecasts.

2 Future value problems

Suppose an investor places $FV_0 =\$1$ in an equity index tracker fund with an investment horizon of H years and wishes to estimate the stochastic distribution of the future value of this fund, FV_H, on this date, or its equivalent annual rate of return $R_H = H^{-1} \ln (FV_H)$.

There are many areas in finance where it is important to understand the properties of FV_H; we briefly note three.[1] First, the expected returns to different asset classes are central inputs into asset allocation decisions; we return to this issue in subsection 4.5 below. Second, there has recently been substantial interest in stochastic asset-liability models. This, in part, has been caused by major swings in the net funding position of defined benefit pension plans.[2] This has led actuaries and pension fund trustees to consider how the funding position of their

[1] The equity premium is a parameter that also has a number of applications beyond investment management. In particular, it is a key input variable in the Capital Asset Pricing Model, which is commonly used in calculating the cost of equity capital for capital budgeting purposes. The techniques that we discuss below would also be relevant in these contexts, but we do not explicitly explore these issues in this paper.

[2] According to the Pension Protection Fund the net funding position of defined benefit pension plans in the UK swung from a surplus of £161.2bn (June 2007) to a deficit of £180.2bn (March 2009), back again to a surplus of

funds might change in the future. Again, the expected return to equity is a critical variable influencing the output of these models. Third, in some countries regulators prescribe the rates which investment providers can use when making illustrations of potential future investment values to retail clients. For example, in the UK, the Financial Conduct Authority (FCA) requires the inclusion of standard deterministic projections for packaged products that are not covered by the European Markets in Financial Instruments Directive (MiFID). The rates to be used are prescribed by the FCA and are justified in documentation that includes detailed reference to the academic literature on the equity premium. (Financial Services Authority, 2012).

Such parties may potentially be interested in two distinct types of estimator for FV_H. First, they may wish to construct a probability density function for FV_H, revealing the mean, standard deviation, skewness and excess kurtosis of future possible asset values. Alternatively, they may want a single "best estimator" of FV_H itself and for the quantile values of FV_H. Here the term "best estimator" might refer either to an unbiased or a minimum least squared error estimator. We consider probability density functions in Section 3 and best estimators in Section 4.

The most common method for estimating the future value of tracker funds is to decompose ex-post equity market returns into three parts. Let r_{mt} denote the single-period realized logarithmic return to equity over the period $[t-1, t]$. We divide $r_{mt} = r_{ft} + \lambda_t + e_t$ where r_{ft} is the compounded single-period risk-free rate, λ_t the compounded single-period ex-ante equity premium, and e_t the single-period ex-post stock market noise. This implies that $FV_H = \exp\left(H\left(r_{fH} + \lambda_H + e_H\right)\right)$ where $r_{fH} = H^{-1}\sum_{t=1}^{H} r_{ft}$, $\lambda_H = H^{-1}\sum_{t=1}^{H}\lambda_t$, $e_H = H^{-1}\sum_{t=1}^{H} e_t$. The future value of the fund, FV_H, will depend on the individual behavior of these three variables and the interaction between them. In order to simplify the exposition throughout this paper and also to focus attention on the equity premium, we assume that r_{fH} is a known scalar at time 0, that λ_H and e_H are independent random variables, and that the statistical process describing e_H is known with certainty.

There are several methods by which the equity premium, λ_H, can be estimated. The traditional approach has been to take an average of long-run historical excess equity returns and assume that a similar average risk premium will hold into the future. This approach has now largely fallen out of favor. Stock index returns are highly noisy and therefore, even if the ex ante equity premium is constant, a very long time period of data is needed to estimate it with any precision. For example, if excess stock returns are identically and independently normally distributed (i.i.n.d.) with an annual standard deviation of 16%, then, even with 100 years of data of any frequency, the 95% confidence interval for the equity premium is approximately the mean estimate ±3%. To improve accuracy, therefore, most studies in the field that are based on historic data go back to the nineteenth century. However, there is increasing evidence that over the past hundred years the equity premium has changed at least once (see, for example, Jagannathan et al., 2000; Fama and French, 2002; Lettau et al., 2008; Freeman, 2011). Identifying when, and by how much, the ex ante equity premium changed introduces additional complexity and increases further the estimator's confidence intervals.

Taking a theoretical approach to estimating the equity premium also has a lengthy tradition. The problem now is that, as shown initially by Mehra and Prescott (1985), the standard model produces a risk premium estimate that is considered by most financial economists to be unrealistically low. Voluminous subsequent literature (see, for example, Mehra and Prescott 2003, Hung and Wang 2005), has failed to produce a new model that has gained widespread acceptance. This approach, therefore, has never been widely used for practical purposes.

More recently many researchers have returned to the problem of predicting time variations in the equity premium. The latest available evidence points towards there being predictability

£50.1bn (January 2011) before returning to a deficit of £317.0bn (May 2012). At the time of writing, the deficit stands at £27.6bn (December 2013).

that is both statistically significant out-of-sample and relevant for asset allocation purposes; Rapach and Zhou (2011) provide a recent review. This has influenced recent work on long-run portfolio construction and asset-liability management problems (e.g., Barberis, 2000; Ferstl and Weissensteiner, 2011). In parallel to this a substantive literature has also grown up reporting the results of surveys on the equity premium (Welch 2000 with updates in 2001, 2007 and 2009; Fernandez and del Campo 2010a, 2010b; the Duke/CFO Magazine surveys most recently reported upon by Graham and Harvey 2010). This approach not only has the advantage of being explicitly forward-looking but also appeals to investors' psychological preferences. Önkal et al. (2009) show in the context of stock market predictions that people place more weight on forecasts made by human experts than those made by statistical models. Despite this, and in contrast to the work on predicting time-variation in the equity premium, no previous research has considered in detail how the data provided by such surveys can be used to address asset allocation and other investment management problems.

The results that we present are based on the data reported in the 2009 update of Welch (2000). We concentrate on the variable $\lambda_i = \ln(1+G_i)$, where G_i is the i^{th} individuals response to the question "I expect the average geometric* equity premium over the next 30 years to be (relative to rolling future contemporaneous short-term (3 month) T-Bills). Note: *The geometric equity premium is 'casual' usage. Think of it as compounded equity return minus compounded risk-free." We choose this response because, as shown by, for example, Jacquier et al. (2005), if stock returns have a constant expected return, μ, the geometric average return $1+G = \exp(\mu)$.[3] This drew n =131 responses.

It is important for the purposes of this paper to distinguish between two types of uncertainty. First, each expert will have his or her own doubts about the true value of the equity premium; Graham and Harvey (2010) present evidence on this. This is not relevant for the analysis contained in this paper and is, therefore, not discussed further. Each expert presents a spot estimate that we interpret as being their best guess as to the true value of the ex-ante equity premium. The second source of uncertainty is that each expert provides a different answer. Our focus here is on describing optimal ways of combining these different spot responses.

Our methods are based on the assumption that the world contains a very large number of experts, E, each of whom has his or her own forecast of the equity premium, λ_i, $i \in [1, E]$. If we could observe all of these estimates, the derived population frequency distribution, $f\left(\lambda_{i|i\in[1,E]}\right)$, would be a gamma distribution, $\Gamma(\alpha, \beta)$, with shape parameter α and rate parameter β. This is given by:

$$f\left(\lambda_{i|i\in[1,E]}\right) \approx \Gamma(\lambda_i; \alpha, \beta) = \frac{\beta^{\alpha}}{\Gamma(\alpha)} \lambda_i^{\alpha-1} \exp(-\beta\lambda_i), \qquad \lambda_i \in (0, \infty) \tag{1}$$

This is an obvious candidate for consideration for characterizing the range of expert opinions since the equity premium must be positive. This choice mirrors Weitzman (2001) who took this distribution for characterizing survey data on the long-term social discount rate in a highly influential paper in environmental economics; see, for example, Jouini et al. (2010), Weitzman (2010) and Freeman and Groom (In Press) for recent extensions of this framework. Through the properties of $\Gamma(\alpha, \beta)$, the population average expert forecast, $\lambda_p = E^{-1} \sum_{i=1}^{E} \lambda_i = \alpha/\beta$.[4]

Through the survey, we access a small fraction of the available opinions; $n \ll E$. However, as n is still a relatively large number in absolute terms, we would expect the sample frequency

[3] The raw data are available at at http://www.ivo-welch.info/academics/equpdate-results2009.htm. Our results are largely insensitive to the choice of survey question used.

[4] Of course, since E is finite, $f(\lambda_i)$ must be a discrete distribution while $\Gamma(a, \beta)$ is continuous. Therefore, strictly, $\lim_{E\to\infty} \lambda_p = \alpha/\beta$. However, we assume that E is sufficiently large for this distinction to not be of relevance.

distribution, $f\left(\lambda_{i|i\in[1,n]}\right)$ to closely resemble the population frequency distribution $f\left(\lambda_{i|i\in[1,E]}\right)$. When we fit the $n = 131$ survey responses to $\Gamma(a, \beta)$, the maximum likelihood estimates of the parameters values are $\widehat{\alpha} = 7.339$, $\widehat{\beta} = 148.53$.[5] In Figure 1, we present a Q-Q plot to demonstrate the goodness-of-fit between $\Gamma\left(\widehat{\alpha}, \widehat{\beta}\right)$ and $f\left(\lambda_{i|i\in[1,n]}\right)$. By inspection, this figure suggests that the fit is good. More formally, we run a Kolmogorov-Smirnov test to quantify the maximum distance between the cumulative distribution functions of this gamma distribution and that of the sample data. The test statistic is 0.1078; which lies below the 5% critical level. We are, therefore, unable to reject at the standard level of statistical significance the null hypothesis that the raw data are gamma distributed.

[Insert Figure 1 around here]

3 The pdf of the future value of equity portfolios

3.1 The method

The problem that the investor faces is that the "true" value of λ_H is unknown. To make informed judgements about this, she turns to the survey data. We denote by $\overline{\lambda} = n^{-1} \sum_{i=1}^{n} \lambda_i$ the sample mean of the survey responses. Using maximum likelihood methods to estimate the parameters of the gamma distribution, $\overline{\lambda}$ is also given by $\overline{\lambda} = \widehat{\alpha}/\widehat{\beta}$; see footnote 5. However, because of the finite nature of the sample, $\widehat{\alpha}$ and $\widehat{\beta}$ will not be identical to the "true" parameters of the gamma distribution, α, β, that describes the population frequency distribution of opinions. The population mean of expert opinions, $\lambda_p = \alpha/\beta$, will generally not equal the mean of sample responses, $\overline{\lambda}$.

Our methods involve using the observed individual responses, $\lambda_1, ..., \lambda_n$, to construct a probability density function, $g_p(\lambda_p)$, for the unobserved value of λ_p. As we note below, there are a number of techniques available for this; we use the one proposed by Kulkarni and Powar (2010). While we assume that the population frequency distribution of expert opinions can be well characterized through $\Gamma(\lambda_i; \alpha, \beta)$, there is no assumption that $g_p(\lambda_p)$ will also be gamma distributed; in fact, it will not be. Instead, the properties of this pdf will need to be determined using numerical methods.

Even if the investor fully knew the population mean estimate, then this would not necessarily reveal the true value λ_H. The difference between the population mean and λ_H is referred to as the "bias", ω, where $\omega = \lambda_p - \lambda_H$. This is motivated by an extensive literature on analysts bias in areas such as earnings forecasts; see, for example, Ho and Tsai (2004) and Ramnath et al. (2008). The investor must then distinguish between what she expects the bias to be, μ_ω, from her uncertainty over the extent of the bias, σ_ω.

Given this, the probability density function, $g_\lambda(\lambda_H)$, that the investor can derive for the true value of the equity premium based on the observed survey data comprises two elements of uncertainty. The first arises from the fact that only 131 experts were asked, whose aggregated views will not fully reflect entire expert opinion even if all responses are drawn from the same underlying statistical distribution. The second reflects the biasedness of analysts forecasts; even if we were to ask an infinite number of experts, we would be unlikely to derive the true value

[5] The maximum likelihood estimators that we use for $\widehat{\alpha}$ and $\widehat{\beta}$ given the sample of responses are reported in, *inter alia*, Grice and Bain (1980): $\ln(\widehat{\alpha}) - \psi(\widehat{\alpha}) = \ln\left(\overline{\lambda}/\widetilde{\lambda}\right)$ and $\widehat{\alpha}/\widehat{\beta} = \overline{\lambda}$. Here $\psi(\cdot)$ is the digamma function and $\overline{\lambda}, \widetilde{\lambda}$ respectively represent the arithmetic and geometric means of the survey data. We use these estimators throughout, even when we assume there is correlation between expert forecasts. In unreported results, available on request from the authors, we demonstrate that our results have low sensitivity to the estimated values of $\widehat{\alpha}, \widehat{\beta}$.

of λ_H. Given this, our methods first consider the distribution $g_p(\lambda_p)$ and then the additional uncertainty caused by biasedness.

There are a number of known numerical techniques (Grice and Bain 1980; Jensen and Kristensen 1991; Wong 1993) for determining the pdf of the population mean of a gamma distribution, $g_p(\lambda_p)$ conditional on observing a finite number of drawings from the distribution. The one that we employ here involves transforming the data into $Y_i = \lambda_i^p$ for some power p (not to be confused with the p subscript, which is used throughout to denote "population"). This adapts the underlying distribution from gamma to Gaussian. Wilson and Hilferty (1931) and Hernandez and Johnson (1980) show that, if the $\lambda_i s$ are gamma distributed, then for $p = 1/3$, the transformed variables Y_i will be approximately normally distributed. Hawkins and Wixley (1986) propose instead $p = 1/4$. In this paper, we follow the Optimal Power Normal Approximation Method (OPNAM) method of Kulkarni and Powar (2010) who propose using $p = 0.246$ when $\alpha > 1.5$ as is the case for Welch's data. When we make the transformation $Y_i = \lambda_i^{0.246}$ we find that the sample mean $\overline{Y} = 0.4712$, the sample standard deviation $\sigma_Y = 4.40\%$, the skewness is -0.52 and the excess kurtosis is 1.67. A Kolmogorov-Smirnov test is unable to reject at a 5% level of significance the null hypothesis that these transformed data are normally distributed.

Now denote by Y_p the population mean of the transformed forecasts; $Y_p = E[\lambda_i^p]$. By the properties of a gamma distribution (see the appendix for a proof),

$$Y_p = E[\lambda_i^p] = E[\lambda_i]^p \frac{\Gamma(\alpha+p)}{\alpha^p \Gamma(\alpha)} = \lambda_p^p \frac{\Gamma(\alpha+p)}{\alpha^p \Gamma(\alpha)} \tag{2}$$

To proceed from here, let $g_Y(Y_p)$ denote the pdf for the value of Y_p that the investor can construct conditional on the sample of transformed responses. Let $u_i = Y_i - Y_p$ be the difference between expert i's transformed response and the transformed population mean. Because of the Gaussian nature of the Y_is, we assume that the vector $\mathbf{u}$ with elements u_i is multivariately normally distributed with zero mean: $\mathbf{u} \sim \mathbf{N}(\mathbf{0}, \mathbf{\Sigma}_u)$ where $\mathbf{0}$ is an $n-$vector of 0s and $\mathbf{\Sigma}_u$ is the variance-covariance matrix of the elements u_i.

We then use the "Unknown Σ" method of Winkler (1981) to construct confidence intervals for Y_p. While we assume that the correlation between expert forecasts is known precisely (we discuss this further below), the variance of their individual forecast errors, $Var[u_i]$, cannot be estimated perfectly. This is because the survey sample size of Welch is relatively small, leaving imprecision in our volatility estimate. Winkler's method relies on observing that when $\mathbf{\Sigma}_u$ is unknown, the pdf of Y_p, $g_Y(Y_p)$, is Student's $t-$distributed with $\delta + n - 1$ degrees of freedom with mean m^* and variance ${s^*}^2$ where:

$$m^* = \frac{\mathbf{1}'\widehat{\mathbf{\Sigma}}_u^{-1}\mathbf{Y}}{\mathbf{1}'\widehat{\mathbf{\Sigma}}_u^{-1}\mathbf{1}}, \qquad {s^*}^2 = \frac{\delta + (m^*\mathbf{1} - \mathbf{Y})' \widehat{\mathbf{\Sigma}}_u^{-1} (m^*\mathbf{1} - \mathbf{Y})}{(\delta + n - 3)\, \mathbf{1}'\widehat{\mathbf{\Sigma}}_u^{-1}\mathbf{1}} \tag{3}$$

δ is a parameter of the Inverse Wishart distribution that is used to describe our Bayesian prior belief about $\mathbf{\Sigma}_u$, $\widehat{\mathbf{\Sigma}}_u$ is the sample estimate of this variance-covariance matrix, $\mathbf{Y}$ is an $n-$vector with elements Y_i and $\mathbf{1}$ is an $n-$vector of 1s. In the appendix we discuss in detail how we parameterize this model based on Welch's sample data.

Conditional on Y_p being Student's $t-$distributed, it is now straightforward to construct L_x, the lower $x/2$ quantile value for $g_Y(Y_p)$. Then, by reversing equation 2, the equivalent quantile value for $g_p(\lambda_p)$, L_x^*, is given by:

$$L_x^* = \left(L_x \hat{\alpha}^p \Gamma(\hat{\alpha}) / \Gamma(\hat{\alpha}+p)\right)^{1/p} \tag{4}$$

where $\hat{\alpha}$ is the maximum likelihood estimator of α. This allows for the numerical construction of $g_p(\lambda_p)$. In the appendix we describe the results of simulations that assess the accuracy of the combined use of the OPNAM method with Winkler's technique for combining correlated forecasts. We find that there is good precision in the estimates L_x^* even with high levels of correlation between expert forecasts.

3.2 Unbiased and independent expert forecasts

We begin by presenting results for the case when there are n independent and unbiased ($\omega = 0$ so $g_p(\lambda_p) = g_\lambda(\lambda_H)$) forecasts of the $H-$period ex-ante equity premium, λ_H, each with identical forecast error variance; $Var[u_i] = \sigma_u^2$ for all i. This means that $\widehat{\Sigma}_{u|ii}^{-1} = \widehat{\sigma}_u^{-2}$ and $\widehat{\Sigma}_{u|ij}^{-1} = 0$ for $j \neq i$, where $\widehat{\Sigma}_{u|ij}^{-1}$ denotes the element in the i^{th} row and j^{th} column of $\widehat{\mathbf{\Sigma}}_u^{-1}$. Under these assumptions, which are relaxed in the next two subsections, $\widehat{\sigma}_u^2$ is estimated by the sample variance of the transformed estimates, $\widehat{\sigma}_Y^2$.[6] Applying equation 3 is, in this case, straightforward and, in Panel A of Table 1, we report the first four moments of $g_\lambda(\lambda_H)$ from the combined application of the Winkler and OPNAM approaches.

[Insert Table 1 around here]

The first column, $n = \infty$, gives the case when there is no uncertainty over the equity premium and therefore $\lambda_H = \widehat{\alpha/\beta}$ with certainty. The second column, $\rho = 0$, presents the case of uncorrelated experts. The moments can be generated here directly from the empirical estimate of $g_\lambda(\lambda_H)$. The presence of uncertainty has slightly increased the mean from 4.941% to 4.947% and is further seen through the non-zero standard deviation. The skewness and excess kurtosis remain close to zero as the number of degrees of freedom of the Student's $t-$distribution, $\delta+n-1$, is over 140 in all cases. This distribution is, therefore, very similar to a normal distribution. We will return to Panel B and the final four columns of Panel A below.

Now consider the future value of an equity portfolio in $H = 30$ years. In order to focus attention on the equity premium it is assumed that both r_{fH} and $\mathbf{\Sigma}_e$, the $H \times H$ variance-covariance matrix of stock market noise, e_t, over the interval $[0, H]$, are known by the investor. The assumption that $\mathbf{\Sigma}_e$ is known perfectly is common in this stream of literature (Blume 1974; Indro and Lee 1997; Cooper 1998; Jacquier et al. 2003, 2005). It is considered reasonable as returns volatility can, in principle, be estimated to within any required confidence interval provided that historic data can be obtained with sufficiently high frequency and returns are i.i.n.d. Results are based on $r_{fH} = 0$ and $\mathbf{1}'\mathbf{\Sigma}_e\mathbf{1} = H\sigma_e^2$ with $\sigma_e = 16\%$, which is broadly consistent with the observed annual volatility of major stock market indices. Results are generated by randomly drawing across 10 million simulations independent values of λ_H, e_H from their respective probability density functions and then calculating $FV_{30} = \exp(30(\lambda_H + e_H))$ for each simulation. We then calculate the mean value of FV_{30} across the simulations and also the 5th, 25th, 50th (median), 75th and 95th quantile values. In Panel A of Table 2 we present equivalent annual rates of return, $R_{30} = 30^{-1}\ln(FV_{30})$, for each of these statistics.

[Insert Table 2 around here]

The first column, $n = \infty$, presents the case when there is no uncertainty over the equity premium. In this case FV_{30} is lognormally distributed and so the mean value and quantiles

[6] It is, in general, necessary to distinguish between the observed cross-sectional variance of the transformed sample data, σ_Y^2, and the variance of u_i, σ_u^2, which reflects the accuarcy of an individual forecast. For example, if all experts make identical, but equally incorrect, forecasts then $\sigma_Y^2 = 0$ (all forecasts are the same) but $\sigma_u^2 > 0$ (all experts are in error). See Freeman and Groom (In Press) and the appendix to this paper for more detailed discussions of this point.

can be calculated directly from the associated normal distribution. The uncertainty over future values is clear with the 90% confidence interval being [0.136%, 9.746%]. Even the 50% confidence interval is wide at [2.971%, 6.911%]. This, of course, understates true uncertainty as the volatility and risk-free rate processes are assumed here to be known with certainty.

The second column, $\rho = 0$, presents the equivalent statistics based on independent and unbiased experts using Welch's data. Our central finding is that the results are highly similar to the $n = \infty$ case, although again there is a slight increase in the mean value from 6.221% to 6.231%. Given the practical difficulties associated with accurately estimating the ex-post stock market noise process, $\mathbf{1}'\mathbf{\Sigma}_e\mathbf{1}$, these differences seem to be of little economic importance. The statistical rationale underlying these results is clear. The standard deviation of survey data on the equity premium is low; 1.77% in the case of Welch's 2009 questionnaire. The central limit theorem implies that, when $n = 131$, the standard error associated with the mean of the distribution is only 15 basis points. If we were to use historic i.i.n.d. market returns with standard deviation of 16% per year then we would need over 10,000 years of data to achieve similar precision. Uncertainty over the true value of the equity premium, λ_H, is swamped by the ex-post stock market noise, e_H.

3.3 Non-independent expert forecasts

Most financial economists read the same academic papers, textbooks and newspapers and live within the same professional environment. There are also a number of psychological and incentive-driven reasons why an individual expert may prefer not to be too far out of line with the overall consensus. It is, therefore, highly likely that any individual forecast will be influenced by the opinions of other experts. As Graham (1996, pp.193–194) notes: "economists have a tendency to 'clump' (make similar forecasts) ... *(this)* simply reveals the high positive correlation in economists' predictions that occurs because they all study the same economic fundamentals and, given similar training, are likely to interpret the inputs in a similar manner. Therefore, clumped forecasts are quite natural; as a result, it may be important to account for correlation when combining more than one forecast". Clemen and Winkler (1985, p.428) also note, in the context of selective judgements from different experts: "The different sources might utilize some common data, share common assumptions, or have access to some of each other's opinions. We can represent such redundance stochastically in terms of positive dependence among the information sources."

This effect of herding, for both informational and non-informational reasons, has been well documented in both the financial and macroeconomic literature; see, for example, Jegadeesh and Kim (2010). Herding has also been observed in recent institutional market trading behaviour (e.g. Singh 2013). If herding also influences equity premium forecasts then the information content contained with n estimates is less than would be revealed by n truly independent forecasts. Fernandez and del Campo (2010a) show that this is likely to be the case as approximately 20% (50%) of their respondents justify their answer by "historic data" ("reference to books or articles").

To model this, following Kotz and Adams (1964), we assume that the transformed estimates, Y_i, are exponentially correlated. If we rank these in ascending order, it is likely that those that are closest together are most influenced by the same schools of thought. Therefore, if we use i, j to denote the rank ordering of two estimates, we model the correlation $Corr\,(Y_i, Y_j) = \rho^{|i-j|}$ for some constant $\rho \in [0, 1)$. We continue to assume that the variance of the individual forecast error, $Var[u_i] = \sigma_u^2$, is the same for all experts. In this case Freeman and Groom (In Press) prove that:[7]

$$\widehat{\mathbf{\Sigma}}_u^{-1} = \kappa\widehat{\sigma}_Y^{-2} \begin{bmatrix} a & b & 0 & ... & 0 \\ b & c & b & ... & 0 \\ 0 & b & c & ... & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & ... & a \end{bmatrix} \tag{5}$$

$$a = \frac{1}{1-\rho^2}, \qquad b = \frac{-\rho}{1-\rho^2}, \qquad c = \frac{1+\rho^2}{1-\rho^2}, \qquad \kappa = \frac{n^2(1-\rho)^2 - n(1-\rho^2) + 2\rho(1-\rho^n)}{n(n-1)(1-\rho)^2}$$

and this can be used in equation 3 as before. The mean is now affected by ρ with:

$$m^* = \overline{Y} + \frac{(n-2)\rho}{n-(n-2)\rho}\left(\overline{Y} - Y_R\right) \tag{6}$$

where $Y_R = (n-2)^{-1}\sum_{i=2}^{n-1} Y_i$, which is the mean of the transformed forecasts after ignoring the highest and lowest responses. This means that for $\rho \neq 0$, $m^* \gtreqless \overline{Y}$ if $Y_R \lesseqgtr \overline{Y}$, with the inequality getting stronger for larger values of ρ. In the case of Welch's data $Y_R = 0.4715 > 0.4712 = \overline{Y}$, resulting in a mean estimate of λ_H that decreases with increased ρ. This can be seen in the final four columns of Panel A of Table 1. Results are presented for the first four moments of $g_\lambda(\lambda_H)$ at a horizon $H = 30$ years based on $\rho \in \{0.25, 0.5, 0.75, 0.9\}$. These can be broadly interpreted as equivalent to cases where a survey had been run with $N \in \{79, 44, 19, 7\}$ independent and unbiased experts respectively.[8] As might be expected, the greater the correlation between experts, the less sure the investor is about the true value of λ_H. Therefore as ρ increases, so the standard deviation, skewness and excess kurtosis of $g_p(\lambda_p)$ all rise. However, these effects are relatively minor in terms of economic significance.

The final four columns of Panel A of Table 2 present the mean and five quantile values of R_{30} when expert opinion is correlated. Again, the results do not differ greatly from the $n = \infty$ case. Even when $\rho = 0.9$, the mean estimate has only fallen from 6.22% to 6.17% while the 90% confidence interval widens from 9.61% to 9.69% compared to $n = \infty$.

3.4 Biased expert forecasts

An additional complication arises from the fact that the population of experts may have a biased forecast of the true ex-ante equity premium; $\lambda_p = \lambda_H + \omega$ for $\omega \neq 0$.

Panel B of Table 1 reports the first four moment of $g_\lambda(\lambda_H)$ when the distribution of the population mean is independent of the bias and $\omega \sim N(\mu_\omega, \sigma_\omega^2)$. For illustration purposes we set $\mu_\omega = -1\%$, $\sigma_\omega = 1\%$ which assumes that analysts have an expected negative bias in estimating the equity premium, possibly as a result of placing too much weight on recent evidence of a declining cost of equity capital. A standard deviation of $\sigma_\omega = 1\%$ implies that the

[7] The term $\widehat{\sigma}_u^2 = \widehat{\sigma}_Y^2/\kappa$ is the estimated variance of each transformed forecast error. However, it is σ_Y^2 that is observed and not σ_u^2, and this is reflected in this formula. When $\rho = 0$, $\widehat{\sigma}_Y^2 = \widehat{\sigma}_u^2$, but this is not true more generally. See also footnote 6 and the appendix.

[8] See Freeman and Groom (In Press) for a discussion on how N is calculated. This is based on the assumption that Σ_u is known rather than estimated with error and we do not alter that analysis here. We make this simplification because N is an indicative number only and is not used in further calculations.

95% confidence interval for the true value of the bias is approximately 4% in spread which we subjectively interpret as being reasonably high uncertainty. The first column, $n = \infty$, presents the case when there is no uncertainty over the population mean survey response. In contrast to Panel A the standard deviation of $g_\lambda(\lambda_H)$ is no longer zero but is instead equal to σ_ω. The influence of the bias remains irrespective of how many experts we ask. However, since ω is normally distributed, there is no skewness or excess kurtosis in $g_\lambda(\lambda_H)$. The second column gives the results for $\rho = 0$. The first four moments can be derived directly from Panel A and the first two moments of ω.[9] All four moments remain very similar to the $n = \infty$ case, although again there is a small increase in the mean estimate from 5.941% to 5.947%. This implies that the effects of the bias dominate the uncertainty over the population mean of opinions on the value of the ex-ante equity premium. When experts are non-independent as well as biased, as given in the final four columns of this panel, results remain similar to the $n = \infty$ case. The increases in standard deviation, skewness and kurtosis caused by correlated forecasts are substantially lower than in the corresponding cases given in Panel A.

Panel A of Table 3 can be interpreted in the same way as Panel A of Table 2 but now there is a bias in the sample responses. For $n = 131$ we again run 10 million simulations and take independent values of λ_p, e_H and ω from their respective distributions to estimate the probability density function of FV_H. In the $n = \infty$ case the mean response is substantially above the equivalent value given in Table 2: 7.371% against 6.221%. There are two reasons for this. First, the mean is shifted upwards by 1% to compensate for the expected bias. Second, the uncertainty in the true level of the bias introduces an Itô effect in the mean, again pushing it up by approximately $0.5H\sigma_\omega^2 = 0.15\%$. For the other columns the results are similar to Panel A of Table 2 when compared against the $n = \infty$ baseline. There is a small increase in the mean for $\rho = 0$ compared to $n = \infty$ (7.383% against 7.371%), but the mean then decreases with increasing ρ. There is also some widening of the 90% confidence interval as ρ increases, but these effects do not seem to be of great practical significance. Therefore, in the case of biased forecasts, understanding the properties of the bias is more important than the nature of the correlation structure.

4 The unbiased and MLSE estimators of future value

4.1 The method

To calculate unbiased and MLSE estimates of FV_H, we adapt the technique of Jacquier et al. (2003, 2005). We take an average of the n estimates given in the sample, $\overline{\lambda} = n^{-1}\sum_{i=1}^{n}\lambda_i$, and choose adjustment factors A_1 and A_2 as follows:

$$\begin{array}{ll} E\left[\exp\left(H\left(\overline{\lambda} - A_1\right)\right) - \exp\left(H\lambda_H\right)\right] = 0 & \text{Unbiasedness} \\ \min\limits_{A_2} E\left[\left(\exp\left(H\left(\overline{\lambda} - A_2\right)\right) - \exp\left(H\lambda_H\right)\right)^2\right] & \text{MLSE} \end{array} \tag{7}$$

This leads to:

$$\begin{aligned} A_1 &= \frac{\ln(I_1)}{H} - \lambda_H \\ A_2 &= \frac{1}{H}\ln\left(\frac{I_2}{I_1}\right) - \lambda_H \end{aligned} \tag{8}$$

[9] When x and y are independent, $\sigma_{(x+y)}^2 = \sigma_x^2 + \sigma_y^2$, $skew(x+y) = \left(\sigma_x^3 skew_x + \sigma_y^3 skew_y\right)/\sigma_{x+y}^3$ and $kurt(x+y) = \left(\sigma_x^4 kurt_x + \sigma_y^4 kurt_y\right)/\sigma_{x+y}^4$.

where $I_1 = E\left(\exp(H\overline{\lambda})\right)$ and $I_2 = E\left(\exp\left(2H\overline{\lambda}\right)\right)$.[10]

Following Kan and Zhou (2009), the adjustor A_1 can also be used to construct the unbiased estimator of the q^{th} quantile of future wealth, FV_q. If λ_H were known then $FV_q = \exp\left(\lambda_H H + z_q\sqrt{\mathbf{1'\Sigma}_e\mathbf{1}}\right)$, where z_q is the q^{th} quantile of a standard normal distribution. If we replace λ_H with the sample mean minus the adjustor, $\overline{\lambda} - A_1$, then

$$E\left[\exp\left((\overline{\lambda} - A_1)H + z_q\sqrt{\mathbf{1'\Sigma}_e\mathbf{1}}\right) - FV_q\right] = 0 \tag{9}$$

which is again unbiased. It is straightforward to extend this argument to show that A_2 is also the MLSE adjustor for the q^{th} quantile.

Notice that the problem is now different to that presented in the previous section. Here the analysis centers on understanding the properties of the sample mean of a survey drawn from a known probability density function, while in the previous section we analyzed the probability density function of the population mean contingent on survey responses. For a normal distribution the pdf of the population mean around the sample mean is the same as the pdf of the sample mean around the population mean. This symmetry, though, does not hold more generally. For this reason, we will need to use an entirely different set of techniques in this section to the previous one.

Our contribution in this section is to extend this framework to gamma distributed survey responses. The adjustors A_1 and A_2 in equation 8 are of the form $E[\exp\left(H\overline{\lambda}\right)]$ and $E[\exp\left(2H\overline{\lambda}\right)]$. These expressions are H^{th} and $2H^{th}$ moment generating functions of the probability density function of $\overline{\lambda}$. One of the main analytical conveniences of the gamma distribution is that its mgfs are well known so that if conditions arise where the sample mean is also gamma distributed, $\overline{\lambda} \sim \Gamma(A, B)$, then:

$$E[\exp\left(x\overline{\lambda}\right)] = \left(1 - \frac{x}{B}\right)^{-A}, \qquad f(\overline{\lambda}) = \Gamma(A, B), \quad -\infty < x < B \tag{10}$$

This equation can be used in conjunction with equation 8 to calculate A_1 and A_2 subject to the respective regularity conditions $B > H$ and $B > 2H$ when $\overline{\lambda} \sim \Gamma(A, B)$:

$$\begin{aligned} A_1 &= -\frac{A}{H}\ln\left(1 - \frac{H}{B}\right) - \lambda_H \\ A_2 &= -\frac{A}{H}\ln\left(1 - \frac{H}{B - H}\right) - \lambda_H \end{aligned} \tag{11}$$

4.2 Unbiased and independent expert forecasts

In the case of independent expert forecasts $\overline{\lambda} = n^{-1}\sum_{i=1}^{n}\lambda_i$ is a constant multiplied by the sum of independently and identically gamma distributed random variables. This means that $\overline{\lambda}$ is also gamma distributed, $\overline{\lambda} \sim \Gamma(n\alpha, n\beta)$. Substituting $A = n\alpha$ and $B = n\beta$ into equation 11 gives the MLSE and unbiased adjustors. While the true parameter values, α, β are unknown, for practical purposes we recommend replacing them with their maximum likelihood estimates, $\widehat{\alpha}, \widehat{\beta}$ and setting $\lambda_H = \widehat{\alpha}/\widehat{\beta}$.

Panels B and C of Table 2 give the equivalent annual rates of return, $R_{30} = 30^{-1}\ln(FV_{30})$, for the unbiased and estimated future value and 5%, 25%, 50%, 75% and 95% quantile values. The

[10] Notice that, if the survey data were unbiased, uncorrelated and normally distributed, $\lambda_i \sim N(\lambda_H, \sigma_\lambda^2)$, then $\overline{\lambda} \sim N\left(\lambda_H, \sigma_\lambda^2/n\right)$. In this case, $I_1 = \exp\left(H\lambda_H + 0.5H^2\sigma_\lambda^2/n\right)$ and $I_2 = \exp\left(2H\lambda_H + 2H^2\sigma_\lambda^2/n\right)$. Given this $A_1 = H\sigma_\lambda^2/2n$ and $A_2 = 3H\sigma_\lambda^2/2n$. These results are reported in Jacquier et al. (2005).

$n = \infty$ case reflects the situation where there is no uncertainty about the true value of the equity premium and therefore all statistics are identical across Panels A, B and C. For the independent errors case, $\rho = 0$, all statistics are very similar to the equivalent statistics for $n = \infty$ and their equivalents in Panel A. This again is a consequence of the low standard deviation of the survey responses and the large number of experts asked, effectively removing 'all' uncertainty about the equity premium value. This contrasts with results presented in Blume (1974), Indro and Lee (1997), Jacquier et al. (2003, 2005) and Kan and Zhou (2009), where the effects of uncertainty over the true expected return derived from historic stock returns are of considerable economic relevance.

4.3 Non-independent expert forecasts

Consider again the case of exponential correlation but now for the untransformed estimates; $Corr(\lambda_i, \lambda_j) = \rho^{|i-j|}$ for ordered survey responses i and j. Kotz and Adams (1964) show that the sum of n observations drawn from this distribution is also approximately gamma distributed with parameters $\Gamma(a_n, \beta_n)$

$$\begin{aligned} \alpha_n &= \left[n + \frac{2\rho}{1-\rho}\left(n - \frac{1-\rho^n}{1-\rho} \right) \right]^{-1} n^2 \alpha \\ \beta_n &= \frac{\alpha_n}{n\alpha}\beta \end{aligned} \tag{12}$$

and therefore $\overline{\lambda} \sim \Gamma(\alpha_n, n\beta_n)$. Notice from the previous subsection (and equation 12 for the case $\rho = 0$) that if we had instead asked N independent and unbiased experts then $\overline{\lambda} \sim \Gamma(N\alpha, N\beta)$. By equating these two expressions, this means that the equivalent number of independent observers is given by $N = \alpha_n/\alpha = n\beta_n/\beta$:

$$N = \left[n + \frac{2\rho}{1-\rho}\left(n - \frac{1-\rho^n}{1-\rho} \right) \right]^{-1} n^2 \tag{13}$$

Setting $\rho \in \{0.25, 0.5, 0.75, 0.9\}$ gives $N \in \{79, 44, 19, 7\}$. These are, to the nearest integer, the same values as given in section 3.3.

Setting $A = N\widehat{\alpha}$ and $B = N\widehat{\beta}$ and substituting these values into equation 11 gives the new adjustment parameters. The final four columns of Panels B and C of Table 2 show how this correlation structure affects the equivalent annual rate of return for the unbiased and MLSE estimators of the future value and quantile values. With high values of ρ there are some clear differences against the $n = \infty$ case with the median future value falling from 4.941% to 4.873% (unbiased) and 4.731% (MLSE). The 90% confidence interval is constant at 9.610% across all examples, which contrasts with the situation in Panel A where the confidence interval widens as ρ grows larger.

An additional finding to note from the first rows of panels A, B and C of Table 2 is that as we introduce uncertainty by moving from $n = \infty$ to $\rho = 0$, so the unbiased and MLSE future values go down while the expected future value taken from the pdf goes up. The former finding is consistent with Blume (1974), Indro and Lee (1997) and Jacquier et al. (2003, 2005) while the latter is analogous to that of Gollier (2004). His results are based on the observation that, by Jensen's inequality, $E[\exp(\widetilde{r}H)] > \exp(E[\widetilde{r}]H)$ when the cost of capital, $\widetilde{r}$, is stochastic. This again emphasizes the difference between taking an unbiasedness approach and a pdf approach.

4.4 Biased expert forecasts

When expert forecasts are biased and potentially correlated, by the law of iterated expectations I_1 and I_2 are given by $I_1 = E\left[E\left(\exp(H\overline{\lambda})|\omega\right)\right]$ and $I_2 = E\left[E\left(\exp\left(2H\overline{\lambda}\right)\right)|\omega\right]$. If $\overline{\lambda} \sim \Gamma\left(N\alpha, N\beta\right)$ we note that $\alpha = \lambda_p^2/\sigma_\lambda^2$ and $\beta = \lambda_p/\sigma_\lambda^2$, where σ_λ^2 is the variance of the population estimates. We will assume in this section that this is exactly equal to the sample variance. Then, as $\lambda_p = \lambda_H + \omega$, from equation 10:

$$E[\exp\left(x\overline{\lambda}\right)|\omega] = \left(1 - \frac{x}{N\left(\lambda_H + \omega\right)/\sigma_\lambda^2}\right)^{-N(\lambda_H+\omega)^2/\sigma_\lambda^2}, \qquad -\infty < x < N\beta \tag{14}$$

Substituting this into equation 8 gives:

$$A_1 = \frac{1}{H}\ln\left(E\left[\left(1 - \frac{H}{N\left(\lambda_H + \omega\right)/\sigma_\lambda^2}\right)^{-N\left((\lambda_H+\omega)^2/\sigma_\lambda^2\right)}\right]\right) - \lambda_H \tag{15}$$

$$\begin{aligned} A_2 = &\frac{1}{H}\ln\left(E\left[\left(1 - \frac{2H}{N\left(\lambda_H + \omega\right)/\sigma_\lambda^2}\right)^{-N\left((\lambda_H+\omega)^2/\sigma_\lambda^2\right)}\right]\right) \\ &-\frac{1}{H}\ln\left(E\left[\left(1 - \frac{H}{N\left(\lambda_H + \omega\right)/\sigma_\lambda^2}\right)^{-N\left((\lambda_H+\omega)^2/\sigma_\lambda^2\right)}\right]\right) - \lambda_H \end{aligned}$$

Panels B and C of Table 3 should be interpreted in the same way as the equivalent panels in Table 2 but now there is a bias in the population estimate. Because λ_H is unobservable, we replace it with $\lambda_H = \hat{\alpha}/\hat{\beta} - \mu_\omega$ for estimation purposes. For the case $n = \infty$, the bias adjustors $A_1 = \mu_\omega + 0.5H\sigma_\omega^2$ and $A_2 = \mu_\omega + 1.5H\sigma_\omega^2$. This means that, in contrast to Table 2, the statistics in the first column vary between panels. These differences are now of some economic significance. For example, the median estimate is 5.941% under the probability density function approach (Panel A), but 5.491% when taking an MLSE estimator (Panel C). Similar results are seen when $\rho = 0$. As ρ increases so the differences between approach start to become of real economic significance. For example, with $\rho = 0.9$, the pdf approach gives an expected future value rate of 7.316% while the MLSE estimator is almost a full percentage point lower at 6.561%. These convert into projected future values of FV_{30} = \$8.98 and \$7.16 respectively based on an initial investment of FV_0 =\$1. Therefore, when there is both herding and bias in the forecasts, it is important to be careful when handling the survey data.

4.5 An application to optimal asset allocation

We now follow Jacquier et al. (2005) by considering how uncertainty over the true value of the equity premium might impact upon the optimal asset allocation for an investor with a power utility function and an investment horizon of H. This investor must choose an optimal weight, W, to place in the equity market portfolio with the remainder being placed in the risk-free asset. Following Merton (1969) we assume that the future value of this mixed equity/Treasury portfolio evolves in continuous time using a simple geometric Brownian motion process:

$$dFV/FV = \left(r_f + W\left(\lambda_H + 0.5\sigma_e^2\right)\right)dt + W\sigma_e dz \tag{16}$$

Then, by Itô's lemma:

$$d\left(\ln FV\right) = \left(r_f + W\left(\lambda_H + 0.5\sigma_e^2\right) - 0.5W^2\sigma_e^2\right)dt + W\sigma_e dz \tag{17}$$

and integrating out and using the fact that $FV_0 = 1$:[11]

$$\ln\left(FV_H\right) = HR_H \sim N\left(\left(r_f + W\lambda_H + 0.5W\left(1-W\right)\sigma_e^2\right)H, W^2\sigma_e^2 H\right) \tag{18}$$

By setting $W = 1$, it is clear that this process is consistent with the discrete time model described above. The optimization is undertaken by maximizing

$$\max_W E\left[U(FV_H)\right] = \max_W E\left[\frac{FV_H^{1-\gamma} - 1}{1-\gamma}\right] = \max_W E\left[\frac{\exp\left(\left(1-\gamma\right)HR_H\right) - 1}{1-\gamma}\right] \tag{19}$$

where $U\left(FV_H\right)$ is the utility of final wealth and γ is the coefficient of relative risk aversion. When $\gamma = 1$, this utility function takes logarithmic form, $U\left(FV_H\right) = \ln\left(FV_H\right)$. This optimization problem for $\gamma \neq 1$ is equivalent to:

$$\max_W \frac{1}{1-\gamma}\exp\left(0.5H\sigma_e^2\left(1-\gamma\right)\left(W - W^2\gamma\right)\right)E\left[\exp(HW(1-\gamma)\lambda_H)\right] \tag{20}$$

Within the assumptions of this paper:

$$\begin{aligned} E\left[\exp(HW(1-\gamma)\lambda_H)\right] &= E\left[\exp(HW(1-\gamma)\left(\lambda_p - \omega\right))\right] \\ &= E\left[\exp(HW(1-\gamma)\lambda_p\right]E\left[\exp(-HW(1-\gamma)\omega\right] \\ &= \left(1 - \frac{HW(1-\gamma)}{N\beta}\right)^{-N\alpha}\exp\left(-HW\left(1-\gamma\right)\mu_\omega + 0.5H^2W^2\left(1-\gamma\right)^2\sigma_\omega^2\right) \end{aligned} \tag{21}$$

resulting in the optimization problem:

$$\begin{aligned} \max_W \frac{1}{1-\gamma}\exp\left(0.5\sigma_e^2 H\left(1-\gamma\right)\left(W - W^2\gamma\right) - HW\left(1-\gamma\right)\mu_\omega + \right. \\ \left. 0.5H^2W^2\left(1-\gamma\right)^2\sigma_\omega^2\right)\left(1 - \frac{HW(1-\gamma)}{N\beta}\right)^{-N\alpha} \end{aligned} \tag{22}$$

Figure 2 presents solutions to equation 22 for $H \in [1, 50]$ based on $\gamma = 4$, $\mu_\omega = -1\%$ and $\sigma_\omega = 1\%$. The "no bias or uncertainty" line assumes that $\lambda_H = \hat{\alpha}/\hat{\beta}$ and therefore there is no allowance for the mean of the bias, the uncertainty of the bias, or the uncertainty of the population mean. This leads to $W = \left(\hat{\alpha}/\hat{\beta} + 0.5\sigma_e^2\right)/\gamma\sigma_e^2$ which is a result originally presented by Merton (1969). This has the lowest initial weight in equity since it ignores the belief that experts are systematically over-pessimistic in their views. The "no uncertainty" line sets $\lambda_H = \hat{\alpha}/\hat{\beta} - \mu_\omega$ but still does not adjust for uncertainty either over the true level of the bias or the population mean. The "n =infinity" line allows for both μ_ω and σ_ω but not uncertainty in the population mean. The closed form solution for the weight in equity is now:[12]

$$W = \frac{\hat{\alpha}/\hat{\beta} - \mu_\omega + 0.5\sigma_e^2}{\gamma\sigma_e^2 - H\left(1-\gamma\right)\sigma_\omega^2} \tag{23}$$

Consistent with Jacquier et al. (2005) the optimal holding in equity declines with the investment horizon.[13] This is caused by uncertainty over the true value of the equity premium. The

[11] As this utility function has constant relative risk aversion, the proportion invested in the risky asset is invariant to the level of initial wealth.

[12] The result of Jacquier et al. (2005) emerges when $N = \infty$, $\mu_\omega = 0$ and $\sigma_\omega^2 = \sigma_e^2/T$.

[13] Although Tsai and Wu (In Press) find an inverted U-shaped weighting for equity with time to retirement in a model with non-constant income growth and time-varying investment opportunities.

"correlation= 0.9" line is based on numerical optimization of the full model, equation 22, with N taken from equation 13 when $\rho = 0.9$. The increased uncertainty over the true value of the equity premium again reduces the optimal holding in equity. At long time horizons, $H = 50$ years, the results are moderately sensitive to the choice of model. The "n =infinity" line, which allows for uncertainty in the level of bias, recommends a holding in equity that is 9% lower than the "no uncertainty" case; 61.5% against 70.5%. The complete model, which also allows for uncertainty in the population mean, is a further 3% lower at 58.4%.

[Insert Figure 2 around here]

5 Designing equity premium surveys

Our empirical results allow us to make some recommendations for the future design of equity premium surveys. Our key finding is that when experts are independent there is a striking similarity between the results obtained either when using the population mean or the sample mean even from a small survey. There seems to be little advantage to be gained from increasing the sample size beyond 10–25 respondents if expert responses are independently drawn from the same distribution. However, many surveys have very large samples; Fernandez and del Campo (2010b) have over two thousand respondents, for example.

There appear to be two potential justifications for requiring such a large sample size. First, we may believe that different opinions are drawn from the same distribution but with very high levels of herding. Alternatively it is possible that there are separate "schools of thought" on the equity premium. It may then be necessary to have a large sample to ensure that all important opinions are solicited in a way that is broadly proportional to their influence in the overall population. If this is the case, then we find it somewhat surprising that current research tends to concentrate only on one particular sub-set of experts in the area; Graham and Harvey (2010) on practitioners and Welch (2001 and updates) from academics, for example. This is despite the fact that Fernandez and del Campo (2010b) show that academics, companies and practitioners tend to give different average expert responses. Taking a more holistic approach to sampling would appear more advantageous than increasing the sample size taken from a broadly homogenous subpopulation.

A potential way to design the survey so as to address issues of non-independence and schools of thought is to add a textual question that asks respondents to briefly explain how they arrived at their answers. We then might reasonably hypothesize that cluster analysis would reveal a relationship between the numerical estimates and the textual responses. Fernandez and del Campo (2010a), who ask the question "Books or articles that I use to support this number", provide graphical evidence that is consistent with this view. They find that those citing Damodaran as a key influencing reference have a noticeably lower average equity premium than those using Morningstar-Ibbotson. A detailed quantitative analysis of the relationship between the empirical and written response would help enable the researcher to gain a more thorough understanding of the nature of the correlation structure that underlies the survey data. In addition the researcher could consider whether the textual responses represent all the main opinions on the equity premium or whether important schools of thought might be missing from, or potentially under-represented in, the sample of respondents. Finally, the textual responses may give the researcher some view on how 'expert' each respondent is, with the potential to remove noise-respondents on the basis of their reply.

While this textual question would help us better understand the relationship between the sample mean and the population mean it would not provide any further understanding of the population bias. As the ex-ante equity premium is not even observable ex-post, the issue of calibrating the bias does not obviously lend itself to direct empirical analysis. We see two possible solutions to this. The first is a purely literature-based approach that would draw inferences from existing studies on biases in macroeconomics and finance. The second is to ask a connected question in an equity premium survey that could be empirically verified ex-post. While both Graham and Harvey (2010) and Welch ask questions about short-term stock market movements, the volatility of realized market returns is so great that identifying the true drift parameter over such a short period of time is impossible. Given this, we would recommend a question about a variable such as GDP growth which has a much lower standard deviation. Assuming there is some correlation between over-optimism (-pessimism) with respect to short-term economic variables and long-term financial ones, this may enable the researcher to understand something of the properties of the equity premium bias.

6 Conclusion

Recent literature has presented the results of several surveys that ask a sample of experts for forecasts of the forward looking equity premium. The summary statistics of these responses have generated considerable interest. To our knowledge, though, no previous work has considered how optimally to combine all the data contained in any one survey. This paper bridges this gap for investors interested in investment management and optimal asset allocation problems.

We focus on two questions; what is the probability density function for the future value of an equity index tracker fund and what are the unbiased and minimum least square error (MLSE) estimates of the future value, and quantiles of the future value, of this fund? The first of these questions revolves around understanding the density function of the true mean of a distribution conditional on the given sample. The second, in contrast, involves understanding the density function of the mean of a random sample drawn from the population.

Following Weitzman (2001) and others we argue that equity premium survey data are well approximated by a gamma distribution. This allows us to use the numerical OPNAM technique of Kulkarni and Powar (2010) to understand the probability density function of the future value. In addition, we derive closed-form solutions for unbiased and MLSE estimates of future value, and quantiles of future value, thus extending the work of Blume (1974), Indro and Lee (1997), Jacquier et al. (2003, 2005) and Kan and Zhou (2009) when forecasts are unbiased. This is then generalized numerically for cases when expert opinions are biased.

By applying these methods to questions of future value and optimal asset allocation our central conclusion is that, unless the sample is small or expert opinions are highly correlated, then uncertainty over the population mean of forecasts is of little practical significance. As a consequence, if analysts are unbiased, working with the sample mean of the survey data alone is sufficiently accurate in most instances. If analysts are biased then the sample mean should be adjusted by the expected level of the bias but in most, if not all, cases further adjustments for uncertainty over the level of bias are small. Given the forward-looking nature of equity premium surveys and the low levels of uncertainty that they contain, we conclude that the benefits from using surveys to address future value problems significantly outweigh their costs when compared to techniques based on historical average returns.

We also make three recommendations for the future design of equity premium surveys. Greater benefit can often be derived from drawing opinions from different groups of experts rather than from increasing the sample size within a particular part of the population. A brief textual question about how the respondents arrived at their answers may help the researcher understand the nature of the inter-dependence between different expert opinions. Finally incorporating a quantitative question about a verifiable short-term economic variable with low ex-post noise may enable the researcher to draw inferences about the nature of equity premium forecast biases.

References

Barberis N (2000) Investing for the long run when returns are predictable. J Financ 60:225–264.

Blume M (1974) Unbiased estimators of long-run expected rates of returns. J Am Stat Assoc 69:634–663.

Clemen RL and Winkler RL (1985) Limits for the precision and value of information from dependent sources. Oper Res 33:427–442.

Cooper I (1996) Arithmetic versus geometric mean estimators: Setting discount rates for capital budgeting. Eur Financ Manag 2:157–167.

Fama EF, French KR (2002) The equity premium. J Financ 57:637–659.

Fernandez P, del Campo J (2010a) Market risk premium used in 2010 by professors: A survey with 1,500 answers. Working paper; http://ssrn.com/abstract=1606563.

Fernandez P, del Campo J (2010b) Market risk premium used in 2010 by analysts and companies: a survey with 2,400 answers. Working paper; http://ssrn.com/abstract=1609563.

Ferstl R, Weissensteiner A (2011) Asset-liability management under time-varying investment opportunities. J Bank Financ 35:182–192.

Financial Services Authority (2012) Rates of return for FSA prescribed projections. London.

Freeman MC (2011) The time-varying equity premium and the S&P500 in the twentieth century. J Financ Res 34:179–215.

Freeman MC, Groom B (In Press) Positively gamma discounting: Combining the opinions of experts on the social discount rate. Econ J.

Hawkins DM, Wixley RAJ (1986) A note on the transformation of Chi-squared variables to normality. Am Stat 40:296–298.

Hernandez F, Johnson RA (1980) The large-sample behavior of transformations to normality. J Am Stat Assoc 75:855–861.

Gollier C (2004) Maximizing the expected net future value as an alternative strategy to gamma discounting. Financ Res Lett 1:85–89.

Graham JR (1996) Is a group of economists better than one? Than none? J Bus 69:193-232.

Graham JR, Harvey CR (2010) The equity risk premium in 2010. Working paper; http://ssrn.com/abstract=16540

Grice JV, Bain LJ (1980) Inferences concerning the mean of the gamma distribution. J Am Stat Assoc 75:929–933.

Ho L-C, Tsay J (2004) Analysts' forecasts of Taiwanese firms' earnings: Some empirical evidence. .Rev Pac Basin Financ Mark Pol 7, 571-597.

Hung M-W, Wang J-Y (2005) Asset prices under prospect theory and habit formation. Rev Pac Basin Financ Mark Pol 8,1–29.

Indro, DC, Lee, WY (1997) Biases in arithmetic and geometric averages as estimates of long-run expected returns and risk premia. Financ Manage 26(Winter):81–90.

Jacquier E, Kane A, Marcus AJ (2003) Geometric or arithmetic mean: A reconsideration. Financ Anal J 59(Nov/Dec):46–53.

Jacquier E, Kane A, Marcus, AJ (2005) Optimal estimation of the risk premium for the long run and asset allocation: a case of compounded estimation risk. J Financ Economet 3:37–55.

Jagannathan R, McGrattan ER, Scherbina A (2000) The declining U.S. equity premium. Federal Reserve Bank of Minneapolis Quarterly Review 24(4):3–19.

Jegadeesh N, Kim W (2010) Do analysts herd? An analysis of recommendations and market reactions. Rev Financ Stud 23:901–937.

Jensen JL, Kristensen LB (1991) Saddlepoint approximations to exact tests and improved likelihood ratio tests for the gamma distribution. Commun Stat-Theor M 20:1515–1532.

Jouini E, Marin J-M, Napp C (2010) Discounting and divergence of opinion. J Econ Theory 145:830–859.

Kan R, Zhou G (2009) What will the likely range of my wealth be? Financ Anal J 65(July/August):68-77.

Kotz S, Adams JW 1964. Distribution of sum of identically distributed exponentially correlated gamma-variables. Ann Math Stat 35:277–283.

Kulkarni HV, Powar, SK, 2010. A new method for interval estimation of the mean of the gamma distribution. Lifetime Data Anal 16:431–447.

Lettau M, Ludvigson SC, Wachter JA (2008) The declining equity premium: What role does macroeconomic risk play? Rev Financ Stud 21, 1653–1687.

Mehra R, Prescott EC (1985) The equity premium: a puzzle. J Monetary Econ 15:145–61.

Mehra R, Prescott EC (2003) The equity premium in retrospect, in Constantinides GM, Harris M, Stulz R (eds) Handbook of the Economics of Finance, Volume 1B: Financial Markets and Asset Pricing (North Holland).

Merton R (1969) Lifetime portfolio selection under uncertainty: The continuous-time case. Rev Econ Stat 51:247–257.

Önkal D, Goodwin P, Thomson M, Gönül S, Pollock A (2009) The relative influence of advice from human experts and statistical methods on forecast adjustments. J Behav Decis Making 22:390–409.

Ramnath S, Rock S, Shane P (2008) The financial analyst forecasting literature: A taxonomy with suggestions for further research. Int J Forecasting 24:34–75.

Rapach DE, Zhou G (2011) Forecasting stock returns. Working paper, Saint Louis University.

Singh V (2013) Did institutions herd during the internet bubble? Rev Quant Financ Acc 41: 513-534.

Tsai H-J, Wu Y (In Press) Optimal portfolio choice with asset return predictability and non-tradable labor income. Rev Quant Financ Acc.

Weitzman ML (2001) Gamma discounting. Am Econ Rev 91: 260–271.

Weitzman ML (2010) Risk-adjusted gamma discounting. J Environ Econ Manag 60:1–13.

Welch I (2000) Views of financial economists on the equity premium and other issues. J Bus 73–4:501–537.

Wilson EB, Hilferty MM (1931) The distribution of Chi-squares. P Natl Acad Sci USA 17:684–688.

Winkler RL (1981) Combining probability distributions from dependent information sources. Manage Sci 27:479-488.

Wong A (1993) A note on inference for the mean parameter of the gamma distribution. Stat Probabil Lett 17:61–66.

7 Appendix

In this appendix, we present further details of the Kulkarni and Powar (2010) OPNAM method for estimating the density function of the mean of a gamma distribution conditional on a finite sample, the Winkler (1981) method for combining correlated expert forecasts, and the application of these two methods to the problem at hand. We also, by simulation, demonstrate the accuracy of these two techniques.

7.1 Transforming the data

The OPNAM method relies on the fact that if the underlying data $\lambda_i \sim \Gamma(\alpha, \beta)$ then $Y_i = \lambda_i^p$ will be approximately normally distributed for $p = 0.246$ provided that $\alpha > 1.5$. The first four moments of Y_i are known in closed form — we derive the first two here and note that the equivalent equations for skewness and kurtosis are given on page 437 of Kulkarni and Powar (2010). The pdf of a gamma distribution is given by:

$$g(\lambda_i) \sim \Gamma(a, \beta) = \frac{\beta^\alpha}{\Gamma(a)} \lambda_i^{\alpha-1} e^{-\beta\lambda_i} \tag{24}$$

So, the expectation of $Y_p = E[\lambda_i^p]$ is given by

$$\begin{aligned} Y_p &= \int_0^\infty \lambda_i^p \frac{\beta^\alpha}{\Gamma(a)} \lambda_i^{\alpha-1} e^{-\beta\lambda_i} d\lambda_i = \int_0^\infty \frac{\beta^\alpha}{\Gamma(a)} \lambda_i^{p+\alpha-1} e^{-\beta\lambda_i} d\lambda_i \\ &= \frac{\Gamma(a+p)}{\beta^p \Gamma(a)} \int_0^\infty \frac{\beta^{p+\alpha}}{\Gamma(a+p)} \lambda_i^{p+\alpha-1} e^{-\beta\lambda_i} d\lambda_i \end{aligned} \tag{25}$$

The transformation in the last line is 'trivial' in the sense that it is just multiplying both the numerator and denominator by $\beta^p \Gamma(a+\beta)$ and then rearranging. The purpose of this trivial transformation is to ensure that the term within the integral is identical to the probability density function of a gamma distribution with parameters $\alpha + p$ and β. Clearly as this is the integral over the whole support of well-defined pdf, it must equal one. Therefore:

$$Y_p = \frac{\Gamma(a+p)}{\beta^p \Gamma(a)} = \left(\frac{\alpha}{\beta}\right)^p \frac{\Gamma(a+p)}{a^p \Gamma(a)} = E[\lambda_i]^p \frac{\Gamma(a+p)}{a^p \Gamma(a)} \tag{26}$$

and this is equation 2 in the body of this paper.

The cross-sectional variance of the forecasts is given by $\sigma_Y^2 = E\left[\lambda_i^{2p}\right] - E^2[\lambda_i^p]$, which follows directly from the preceding argument:

$$\sigma_Y^2 = \left(\frac{1}{\beta^p \Gamma(\alpha)}\right)^2 \left[\Gamma(\alpha+2p)\Gamma(a) - \Gamma^2(\alpha+p)\right] \tag{27}$$

Similar expressions follow for the skewness and excess kurtosis and Kulkarni and Powar note that these depend on α and p but not β. Therefore the choice of p depends on α only. For $\alpha > 1.5$, they select $p = 0.246$ so as to get the skewness and excess kurtosis jointly close to zero. Their method is an approximation, though, and therefore the $Y_i s$ are only approximately normally distributed even if the $\lambda_i s$ are perfectly gamma distributed.

A key feature to notice is that this transformation from gamma to normal does not depend on the generating process for the underlying $\lambda_i s$. Therefore this applies equally well to correlated as uncorrelated forecasts. To explicitly demonstrate this, we simulated correlated gamma distributed random variables as follows. First, using a Cholesky decomposition approach, we generate

131 normally distributed $N(0,1)$ and exponentially correlated random variables. We then convert these to correlated uniform random variables using the cumulative density function of the standard normal distribution. Finally, we take an Inverse Gamma function to convert again into correlated gamma distributed random variables, λ_i.[14] We then construct the transformations $Y_i = \lambda_i^{0.246}$ and calculate the first four moments of the transformed forecasts.

In Table A, we report our results. In Panel A we present information on the initial correlation structure that is put into the Cholesky decomposition and the estimated correlations between λ_i and λ_{i+1} and Y_i and Y_{i+1} across the simulations (where $i = 50$ is chosen at random for illustrative purposes). This clearly demonstrates that the correlation structure remains almost entirely unaffected by the transformations that we undertake — we discuss this point further below. In Panel B we present the median values of the first four moments of the Y_is across these simulations, along with lower 2.5% and upper 97.5% simulated values. These are compared against the theoretical values given above and in Kulkarni and Powar (2010).

[Include Table A around here]

For all values of ρ considered, the theoretical predictions lie within the 95% confidence interval and close to the observed median values. This shows the robustness of the power transformation to correlation in the experts' forecasts.

7.2 Constructing confidence intervals for α/β

The purpose of using the OPNAM method is to construct confidence intervals, L_x^*, for the population mean (rather than survey mean), λ_p, of expert responses. To do this, we first construct confidence intervals, L_x for the population mean, Y_p (as given in equation 26) of the transformed responses and then use equation 4 to construct L_x^*. We do this using the "Unknown Σ" method described in Winkler (1981). While we assume that the correlation between experts, $\rho^{|i-j|}$, is known perfectly, the sample cross-sectional variance of transformed forecasts, $\widehat{\sigma}_Y^2$, will differ from the population cross-sectional variance of transformed forecasts, σ_Y^2. As $\Sigma_{u|ij} = \sigma_Y^2 \rho^{|i-j|}$, this introduces sampling error into our estimate $\widehat{\Sigma}_u$.

To overcome this problem, following Winkler (1981), we form a Bayesian prior that the "true" variance-covariance matrix is Inverse Wishart distributed with parameter value δ:

$$f\left(\mathbf{\Sigma}|\widehat{\mathbf{\Sigma}},\delta\right) \propto \left|\mathbf{\Sigma}^{-1}\right|^{\frac{\delta+2n}{2}} \exp\left(-\frac{\delta}{2}tr\left(\mathbf{\Sigma}^{-1}\widehat{\mathbf{\Sigma}}\right)\right) \tag{28}$$

where $|\cdot|$ is a matrix determinant and $tr(\cdot)$ is a matrix trace. In this case, as discussed in the body of the paper, Y_p is Student's $t-$distributed conditional on the transformed sample data, with $\delta + n - 1$ degrees of freedom and mean m^* and variance ${s^*}^2$, as given in equation 3. In the case of exponential correlation, $\widehat{\mathbf{\Sigma}}_u^{-1}$ is given in equation 5 and, from here, the expressions for

[14] As discussed in early footnotes, it is necessary to distinguish between the cross-sectional variation of responses, σ_Y^2, and the variance of individual forecast errors, σ_u^2; $\sigma_u^2 = \sigma_Y^2/\kappa$ where $\kappa = 1$ when $\rho = 0$. To account for this distinction, the Inverse Gamma function is calibrated to α^* and β^*, where $\alpha^* = \alpha\kappa^{-1}$ and $\beta^* = \beta\kappa^{-1}$. The formula for κ in the case of exponential correlation is given in equation 5. We discuss this distinction in more detail below.

m^* and ${s^*}^2$ can be simplified by noting that:

$$\mathbf{1}'\widehat{\mathbf{\Sigma}}_u^{-1}\mathbf{1} = \kappa\widehat{\sigma}_Y^{-2}\frac{n(1-\rho)+2\rho}{1+\rho}, \qquad \mathbf{Y}'\widehat{\mathbf{\Sigma}}_u^{-1}\mathbf{1} = \kappa\widehat{\sigma}_Y^{-2}\frac{n\overline{Y}-\rho\,(n-2)\,Y_R}{1+\rho}$$
$$\mathbf{Y}'\widehat{\mathbf{\Sigma}}\mathbf{Y} = \kappa\widehat{\sigma}_u^{-2}\frac{nV+\rho^2\,(n-2)\,V_R-2\rho\,(n-1)\,Z}{1-\rho^2} \tag{29}$$
$$V=\frac{1}{n}\sum_{i=1}^{n}Y_i^2, \qquad V_R=\frac{1}{n-2}\sum_{i=2}^{n-1}Y_i^2, \qquad Z=\frac{1}{n-1}\sum_{i=1}^{n-1}Y_iY_{i+1}$$

Therefore the only complexity that arises when operationalizing equation 3 is identifying the parameter δ.

7.3 Estimating δ

To estimate δ, we draw upon the commonly observed parallel between the Inverse Wishart distribution and the Inverse Gamma distribution. We start with the standard assumption that the relationship between the cross-sectional sample variance and cross-sectional population variance is given by a Chi-squared distribution with $v-1$ degrees of freedom; $(v-1)\left(\widehat{\sigma}_Y^2/\sigma_Y^2\right) \sim \chi_{v-1}^2$. We explain below how v is estimated. This means that:

$$\frac{\widehat{\sigma}_Y^2}{\sigma_Y^2} \sim \frac{\chi_{v-1}^2}{v-1} = \Gamma\left(\frac{v-1}{2}, \frac{v-1}{2}\right) \tag{30}$$

which is a gamma distribution with both shape and rate parameters equal to $(v-1)/2$. It then follows that $\sigma_Y^2 \sim \Gamma^{-1}\left(\frac{v-1}{2}, \frac{v-1}{2}\widehat{\sigma}_Y^2\right)$, where $\Gamma^{-1}(\cdot,\cdot)$ is an inverse gamma function and the second parameter is now the scale parameter. Then, by the properties of the Inverse Gamma distribution:

$$Var\left[\sigma_Y^2\right] = \frac{2(v-1)\widehat{\sigma}_Y^2}{(v-3)^2\,(v-5)} \tag{31}$$

To link this to the Inverse Wishart distribution that forms the basis for the "Unknown Σ" method of Winker, we note $\sigma_Y^2 = \Sigma_{u|ii}$, and therefore the variance of this variable is given by $Var\left[\Sigma_{u|ii}\right]$. For an Inverse Wishart distribution, the relationship between the precision of our estimate of $\Sigma_{u|ii}$ and $\widehat{\sigma}_Y^2$ is given by:

$$Var\left[\Sigma_{u|ii}\right] = \frac{2\delta\widehat{\sigma}_Y^2}{(\delta-2)^2\,(\delta-4)} \tag{32}$$

and by comparison of equations 31 and 32 it is clear that $\delta = v-1$.

To estimate v we take a numerical approach. We run 100,000 simulations for all $\rho \in \{0, 0.25, 0.5, 0.75, 0.9\}$. In each case, we draw 131 exponentially correlated Normal random error terms $u_i \sim N\left(0, \sigma_Y^2\right)$, where for simulation purposes, $\sigma_Y^2 = \widehat{\sigma}_Y^2$. For each simulation, s, we then calculate the cross-sectional variance of the error term, $\widehat{\sigma}_{sY}^2$. We then take the maximum likelihood estimator of $(v-1)/2$, the shape parameter of the gamma distribution, across the 100,000 simulations.[15] The values that emerge for v for $\rho \in \{0,\ 0.25,\ 0.5,\ 0.75, 0.9\}$ are, to the nearest integer, $v = \{131,\ 115,\ 81,\ 40,\ 17\}$. These results are reported in more detail in Table B. This

[15] See, for example, Grice, J.V. and J.T.Bain, 1980, Inferences Concerning the Mean of the Gamma Distribution, Journal of the American Statistical Association, Vol. 75, No. 372, pp. 929-933

table also compares the first four moments of $\widehat{\sigma}_Y^2/\sigma_Y^2$ against those of a $\Gamma\left(\left(v-1\right)/2,\left(v-1\right)/2\right)$ distribution.

[Include Table B around here]

7.4 Accuracy of the method

To demonstrate the accuracy of our joint use of the OPNAM method and Winkler's "Unknown Σ" technique for combining correlated forecasts, we run a further set of simulations. Within each of 100,000 simulations, for each value of $\rho \in \{0, 0.25, 0.5, 0.75, 0.9\}$ we draw 131 values of Y_i from an exponentially correlated normal distribution $N\left(E[Y_i], Var\left[Y_i\right]\right)$. As discussed above, both $E\left[Y_i\right]$ and $Var\left[Y_i\right]$ are known in closed form as functions of α, β. For simulation purposes, we use the values of α and β that are empirically estimated from Welch's data. For each simulation, we use a combination of the Winkler (1981) "Unknown Σ" method described above with the Kulkarni and Powar method to derive one-sided upper and lower 0.5%, 1%, 1.5%,..., 9.5% and 10% confidence intervals for α/β. We then count the proportion of simulations where the initial value of α/β lies outside the estimated confidence interval. These results are reported in Table C.

[Include Table C around here]

As can be see, in all cases there is very close agreement between the estimated confidence level and the proportion of simulations where α/β lies outside the estimated confidence interval. This demonstrates the robustness of the method that we use to correlation in the forecasts.

7.5 The correlation structure

Finally, we turn to the correlation structure and demonstrate again that there is close agreement between the assumed correlation structure for the untransformed and transformed estimates (see also Panel A of Table A above). For the case of $\rho = 0.9$, we simulate 131 correlated $Y_i s$ using the standard Cholesky decomposition approach and then construct values of $\lambda_i = Y_i^{1/0.246}$. We repeat these simulations 10,000 times. In Figure 1 we present the correlation between λ_1 and λ_j for $j \in [2, 131]$ across the 10,000 simulations. This is compared against $\text{Corr}(Y_1, Y_j)$ and $\rho^{|i-j|}$, which is the theoretical value under exponential correlation.

[Include Figure A around here]

It is clear that the correlation structure of the $\lambda_i s$ is also very close to being exponentially correlated with $\rho = 0.9$ even though we are modelling errors on the transformed values Y_i.

Table 1: The first four moments of the pdf of the equity premium

	$n = \infty$	$\rho = 0$	$\rho = 0.25$	$\rho = 0.5$	$\rho = 0.75$	$\rho = 0.9$
			Panel A: No bias			
Mean	4.941%	4.947%	4.944%	4.937%	4.918%	4.865%
Standard Deviation	0.00%	0.17%	0.19%	0.22%	0.28%	0.37%
Skewness	0.000	0.075	0.088	0.103	0.128	0.176
Excess Kurtosis	0.000	0.012	0.015	0.027	0.042	0.069
			Panel B: With bias			
Mean	5.941%	5.947%	5.944%	5.937%	5.918%	5.865%
Standard Deviation	1.00%	1.01%	1.02%	1.03%	1.04%	1.07%
Skewness	0.000	0.001	0.001	0.000	0.004	0.008
Excess Kurtosis	0.000	0.000	-0.001	-0.002	-0.002	0.001

This table describes the first four moments of the pdf of the equity premium $g_\lambda(\lambda_H)$. The horizon is 30 years. The model is calibrated using the 2009 updated data of Welch (2000) and the assumption that experts are unbiased in Panel A and biased in Panel B with a mean bias of -1% and a bias standard deviation of 1%. The risk-free rate is assumed to be 0% and ex-post annual stock market volatility is 16%. The $n = \infty$ case is the situation where there is no uncertainty over the population mean estimate of the ex-ante equity premium. The $\rho = 0$ case presents results where there are n = 131 independent experts. The different values of ρ in the other columns correspond to the degree of correlation between the experts' forecast errors.

Table 2: Mean and quantile estimates of future value returns at $H = 30$ years with unbiased experts

	$n = \infty$	$\rho = 0$	$\rho = 0.25$	$\rho = 0.5$	$\rho = 0.75$	$\rho = 0.9$
	Panel A: From the probability density function					
Mean	6.221%	6.231%	6.229%	6.227%	6.209%	6.165%
5% quantile	0.136%	0.133%	0.128%	0.116%	0.092%	0.021%
25% quantile	2.971%	2.975%	2.970%	2.964%	2.939%	2.877%
50% quantile (Median)	4.941%	4.947%	4.946%	4.938%	4.919%	4.864%
75% quantile	6.911%	6.921%	6.920%	6.916%	6.898%	6.850%
95% quantile	9.746%	9.759%	9.756%	9.757%	9.744%	9.709%
	Panel B: Unbiased estimators					
Expected Future Value	6.221%	6.217%	6.215%	6.210%	6.195%	6.153%
5% quantile	0.136%	0.132%	0.130%	0.125%	0.110%	0.068%
25% quantile	2.971%	2.967%	2.964%	2.959%	2.945%	2.902%
50% quantile (Median)	4.941%	4.937%	4.935%	4.930%	4.915%	4.873%
75% quantile	6.911%	6.908%	6.905%	6.900%	6.885%	6.843%
95% quantile	9.746%	9.742%	9.740%	9.735%	9.720%	9.678%
	Panel C: MLSE estimators					
Expected Future Value	6.221%	6.210%	6.202%	6.187%	6.142%	6.011%
5% quantile	0.136%	0.125%	0.117%	0.102%	0.057%	-0.074%
25% quantile	2.971%	2.959%	2.952%	2.937%	2.892%	2.760%
50% quantile (Median)	4.941%	4.930%	4.922%	4.907%	4.862%	4.731%
75% quantile	6.911%	6.900%	6.892%	6.877%	6.832%	6.701%
95% quantile	9.746%	9.735%	9.727%	9.712%	9.667%	9.536%

This table presents information related to the mean, median and four other quantile future value estimates, FV_{30}, at an investment horizon of $H = 30$ years when expert opinions are unbiased. The table presents equivalent annual rates of return, $R_{30} = 30^{-1} ln(FV_{30})$, for each of the future values. Panel A presents estimates from the probability density function, Panel B presents unbiased estimates and Panel C presents MLSE estimates. The $n = \infty$ case is the situation where there is no uncertainty over the true value of the equity premium. The $\rho = 0$ case presents results where there are $n = 131$ independent and unbiased experts. The different values of ρ in the other columns correspond to the degree of correlation between the experts' forecast errors.

Table 3: Mean and quantile estimates of future value returns at $H = 30$ years with biased experts

	$n = \infty$	$\rho = 0$	$\rho = 0.25$	$\rho = 0.5$	$\rho = 0.75$	$\rho = 0.9$
	Panel A: From the probability density function					
Mean	7.371%	7.383%	7.380%	7.374%	7.360%	7.316%
5% quantile	0.862%	0.863%	0.857%	0.846%	0.820%	0.750%
25% quantile	3.858%	3.862%	3.859%	3.850%	3.826%	3.767%
50% quantile (Median)	5.941%	5.949%	5.945%	5.937%	5.920%	5.865%
75% quantile	8.024%	8.033%	8.032%	8.025%	8.010%	7.963%
95% quantile	11.020%	11.034%	11.032%	11.029%	11.018%	10.981%
	Panel B: Unbiased estimators					
Expected Future Value	7.071%	7.067%	7.065%	7.060%	7.045%	7.002%
5% quantile	0.986%	0.983%	0.980%	0.975%	0.960%	0.917%
25% quantile	3.821%	3.817%	3.815%	3.809%	3.795%	3.752%
50% quantile (Median)	5.791%	5.787%	5.785%	5.780%	5.765%	5.722%
75% quantile	7.761%	7.758%	7.755%	7.750%	7.735%	7.693%
95% quantile	10.596%	10.592%	10.590%	10.585%	10.570%	10.527%
	Panel C: MLSE estimators					
Expected Future Value	6.771%	6.760%	6.752%	6.737%	6.692%	6.561%
5% quantile	0.686%	0.675%	0.667%	0.652%	0.607%	0.476%
25% quantile	3.521%	3.509%	3.502%	3.487%	3.442%	3.311%
50% quantile (Median)	5.491%	5.480%	5.472%	5.457%	5.412%	5.281%
75% quantile	7.461%	7.450%	7.443%	7.427%	7.382%	7.251%
95% quantile	10.296%	10.285%	10.277%	10.262%	10.217%	10.086%

This table is similar to Table 2, except now there is bias in the experts' forecasts. The mean value of the bias is -1% and the standard deviation of the bias is 1%. This means that, even in the $n = \infty$ case, uncertainty remains over the true value of the equity premium.

Table A: The first four moments of the transformed forecasts, Y_i

Panel A: The correlation structure						
Input		$\rho = 0$	$\rho = 0.25$	$\rho = 0.5$	$\rho = 0.75$	$\rho = 0.9$
λ_i, λ_{i+1}		$\rho = -0.0001$	$\rho = 0.2423$	$\rho = 0.4960$	$\rho = 0.7347$	$\rho = 0.8921$
Y_i, Y_{i+1}		$\rho = -0.0005$	$\rho = 0.2515$	$\rho = 0.5039$	$\rho = 0.7437$	$\rho = 0.8995$
Panel B: The first four moments of the Y_is						
	Theory	Simulations				
Mean	0.4711	0.4712	0.4711	0.4709	0.4710	0.4704
		[0.4634, 0.4786]	[0.4616, 0.4808]	[0.4583, 0.4840]	[0.4508, 0.4903]	[0.4370, 0.5038]
Standard Deviation	0.0436	0.0435	0.0435	0.0434	0.0431	0.0422
		[0.0383, 0.0488]	[0.0380, 0.0492]	[0.0370, 0.0506]	[0.0343, 0.0539]	[0.0299, 0.0601]
Skewness	-0.0909	-0.0891	-0.0835	-0.0818	-0.0796	-0.0650
		[-0.5154, 0.3237]	[-0.5033, 0.3291]	[-0.5404, 0.3613]	[-0.6571, 0.4931]	[-0.8153, 0.6391]
Excess Kurtosis	-0.0040	-0.1103	-0.1137	-0.1439	-0.2102	-0.3662
		[-0.6613, 0.9074]	[-0.6600, 0.9285]	[-0.6826, 0.9121]	[-0.7985, 0.9090]	[-1.0130, 0.9729]

We run 10,000 simulations. In each 131 correlated gamma distributed random variables, λ_i are generated. We then construct the transformed variables $Y_i = \lambda_i^{0.246}$ and calculate the first four moments of Y_i within each simulation. In Panel A we present the inputted correlation into the simulations and the estimated values of $Corr(\lambda_i, \lambda_{i+1})$ and $Corr(Y_i, Y_{i+1})$ across the 10,000 simulations (where $i = 50$ is chosen at random). This demonstrates the similarity between the correlation structure on the original and transformed estimates. Panel B then presents the median values of these moments, together with lower 2.5% and upper 97.5% confidence intervals, across the 10,000 simulations. The theoretical figures are constructed in closed form using equations either in this appendix or Kulkarni and Powar (2010).

Table B: Estimates of the degrees of freedom of the Chi-squared distribution

	$\rho = 0$	$\rho = 0.25$	$\rho = 0.5$	$\rho = 0.75$	$\rho = 0.9$
			Degrees of freedom		
MLE estimate of v	131.54	116.83	80.98	39.66	16.94
			Moments		
Observed mean	1.0003	1.0004	1.0004	1.0005	1.0011
Theoretical mean	1.0000	1.0000	1.0000	1.0000	1.0000
Observed standard deviation	0.1239	0.1317	0.1590	0.2313	0.3709
Theoretical standard deviation	0.1238	0.1314	0.1581	0.2275	0.3542
Observed skewness	0.2483	0.2935	0.4146	0.6743	1.1395
Theoretical skewness	0.2476	0.2628	0.3163	0.4549	0.7084
Observed excess kurtosis	0.0969	0.1127	0.2490	0.7128	2.1294
Theoretical excess kurtosis	0.0919	0.1036	0.1500	0.3104	0.7528

For all values of $\rho \in \{0, 0.25, 0.5, 0.75, 0.9\}$ we run 100,000 simulations. Within each, we generate 131 exponentially normally distributed error terms $u_i \sim N(0, \sigma_u^2)$, where the value of σ_u^2 is chosen so that the expected cross-sectional variance is the same as that observed in Welch's data. We then calculate the ratio of the sample cross-sectional variance, $\hat{\sigma}_Y^2$ over the true value of σ_Y^2 for each simulation. This ratio is approximately gamma distributed. Panel B reports the estimated first four moments of the ratio against the theoretical values as given by a gamma distribution based on the maximum likelihood estimator of the shape parameter. This is then used to estimate v, where $v - 1$ is the degrees of freedom for the χ^2-distribution, which is reported in Panel A.

Table C: Examining the confidence intervals of $g_p(\lambda_p)$

One-sided	$\rho = 0$		$\rho = 0.25$		$\rho = 0.5$		$\rho = 0.75$		$\rho = 0.9$	
C.I.	Above	Below	Above	Below	Above	Below	Above	Below	Above	Below
0.50%	0.50%	0.46%	0.50%	0.46%	0.49%	0.45%	0.47%	0.44%	0.46%	0.45%
1.00%	0.99%	0.93%	0.97%	0.93%	0.96%	0.90%	0.95%	0.91%	0.93%	0.89%
1.50%	1.48%	1.39%	1.48%	1.40%	1.47%	1.39%	1.45%	1.37%	1.44%	1.37%
2.00%	1.98%	1.87%	1.96%	1.86%	1.95%	1.89%	1.93%	1.88%	1.92%	1.83%
2.50%	2.44%	2.38%	2.45%	2.36%	2.43%	2.33%	2.44%	2.35%	2.42%	2.32%
3.00%	2.95%	2.85%	2.95%	2.84%	2.95%	2.83%	2.90%	2.82%	2.94%	2.79%
3.50%	3.44%	3.35%	3.43%	3.35%	3.42%	3.33%	3.41%	3.31%	3.43%	3.26%
4.00%	3.95%	3.84%	3.94%	3.84%	3.93%	3.84%	3.92%	3.77%	3.95%	3.72%
4.50%	4.47%	4.31%	4.48%	4.31%	4.47%	4.29%	4.45%	4.24%	4.42%	4.22%
5.00%	4.98%	4.81%	4.99%	4.78%	4.97%	4.73%	4.94%	4.73%	4.91%	4.70%
5.50%	5.48%	5.27%	5.48%	5.27%	5.48%	5.25%	5.44%	5.20%	5.40%	5.20%
6.00%	5.97%	5.73%	5.95%	5.73%	5.95%	5.76%	5.94%	5.69%	5.87%	5.67%
6.50%	6.43%	6.20%	6.44%	6.22%	6.44%	6.20%	6.42%	6.19%	6.34%	6.15%
7.00%	6.92%	6.68%	6.90%	6.66%	6.93%	6.67%	6.91%	6.65%	6.83%	6.64%
7.50%	7.43%	7.20%	7.41%	7.19%	7.38%	7.15%	7.36%	7.13%	7.33%	7.15%
8.00%	7.92%	7.70%	7.90%	7.68%	7.85%	7.64%	7.85%	7.65%	7.82%	7.67%
8.50%	8.38%	8.18%	8.35%	8.16%	8.32%	8.14%	8.33%	8.16%	8.31%	8.19%
9.00%	8.85%	8.67%	8.85%	8.67%	8.83%	8.66%	8.80%	8.67%	8.79%	8.69%
9.50%	9.36%	9.16%	9.34%	9.16%	9.34%	9.16%	9.29%	9.15%	9.29%	9.19%
10.00%	9.82%	9.65%	9.83%	9.65%	9.84%	9.67%	9.80%	9.65%	9.79%	9.66%

For all values of $\rho \in \{0, 0.25, 0.5, 0.75, 0.9\}$ we run 250,000 simulations. Within each, 131 exponentially correlated transformed estimates, $Y_i \sim N(E[Y_i], \sigma^2[Y_i])$ are created, where both $E[Y_i]$ and $\sigma^2[Y_i]$ are functions of α and β, as shown in this appendix. The OPNAM method of Kulkarni and Powar (2010), and the "Unknown Σ" method of Winkler (1981) are then used to create 0.5%, 1%, 1.5%, ... , 10% one-sided upper and lower confidence intervals (C.I.) for the value of α/β. The table presents the proportion of simulations where the true value of α/β lies outside the estimated confidence interval.

Figure 1: Q-Q plot of Welch's 2009 data against a gamma distribution

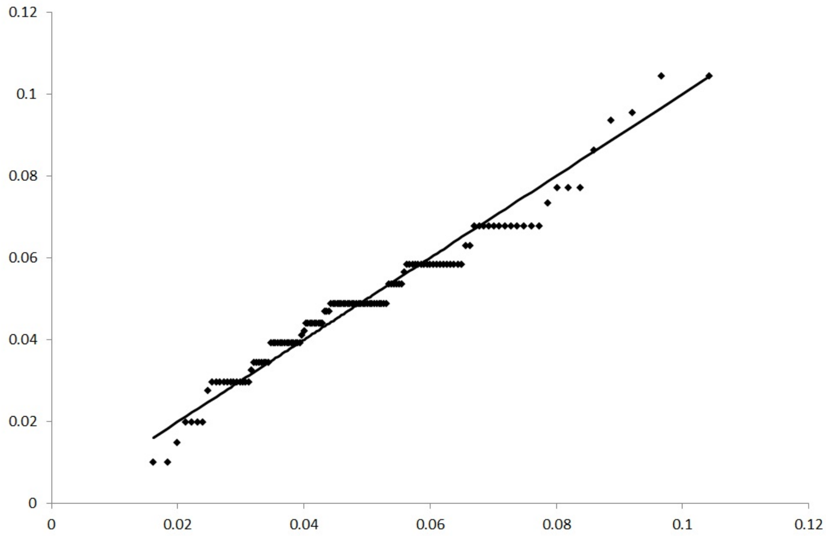

This figure gives a Q-Q plot of Welch's 2009 data against a gamma distribution with parameters estimated by maximum likelihood based on the same data.

Figure 2: Optimal asset allocation in equity

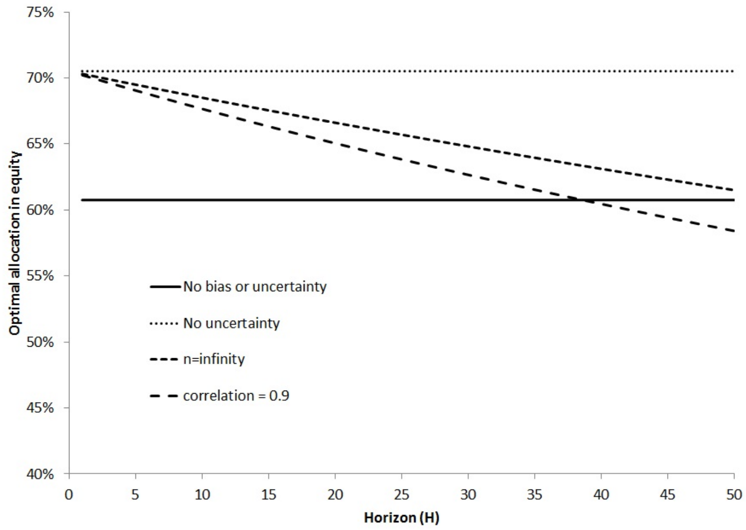

This figure gives the optimal asset allocation in equity based on $\gamma = 4$, $\mu_\omega = -1\%$ and $\sigma_\omega = 1\%$. The "no bias or uncertainty" line treats the sample average as the true value of λ_H and the optimal allocation in equity is given by Merton (1969). The "no uncertainty" line treats $\lambda_H = \alpha/\beta - \mu_\omega$ using the sample estimates of α and β. The "$n =$ infinity" line assumes that $\lambda_p = \alpha/\beta$ using the sample estimates of α and β but now allows for both the mean and uncertainty of the bias. The full model is given by the "correlation =0.9" line, which allows for uncertainty in both the population mean and the bias. This is based upon an exponential correlation structure with $\rho = 0.9$.

Figure A: The correlation between raw forecasts, $Corr(\lambda_1, \lambda_j)$

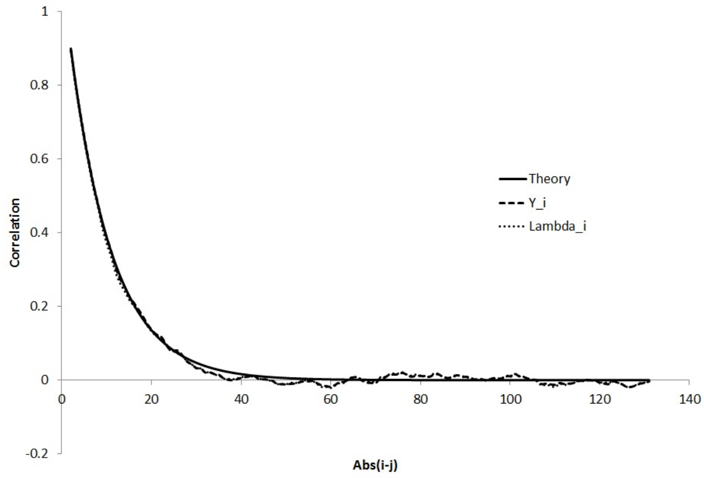

From 131 exponentially correlated normally distributed random variables, Y_i, we estimate the correlation of $\lambda_i = Y_i^{1/0.246}$ against $\lambda_1 = Y_1^{1/0.246}$ for all $i \in [2, 131]$ across $10,000$ simulations. These correlations are then compared against the theoretical value $\rho^{|i-1|}$ and the correlation of Y_i with Y_1. This graph is constructed with $\rho = 0.9$.