Incorporating Economic Objectives into Bayesian Priors: Portfolio Choice Under Parameter Uncertainty

Jun Tu and Guofu Zhou*

Current version: January, 2006

*Singapore Management University and Washington University in St. Louis. We are grateful to Yongmiao Hong and seminar participants at Tsinghua University, Washington University, and the 2005 Finance Summer Camp of Singapore Management University for helpful comments. Correspondence: Guofu Zhou, Olin School of Business, Washington University, St. Louis, MO 63130. Phone: (314) 935-6384 and e-mail: zhou@wustl.edu

Incorporating Economic Objectives into Bayesian Priors:
Portfolio Choice Under Parameter Uncertainty

Bayesian priors on model parameters often ignore the economic objectives at hand. This paper shows that this can be suboptimal, and proposes a way to allow priors to reflect the objective of maximizing an expected utility. Using monthly returns of the Fama-French 25 assets and their three factors from January 1965 to December 2004, we find that the objective-based priors out-perform alternative priors substantially, with annual certainty-equivalent gains of over 10% in many cases. The better performance is present even in repeated sampling experiments, suggesting that objective-based Bayesian optimal portfolios are superior decision rules even judged by the classical statistical criterion.

1 Introduction

Many finance problems have well-defined economic objectives, but model estimation and testing usually make no connection between the estimation of a model to its economic uses. In the classic framework, it is known that different loss functions call for different parameter estimates. Lehmann and Casella (1998) summarize many of the exiting approaches for obtaining such estimates. In contrast, to our knowledge, there are no general theory in a Bayesian set-up to guide one to form priors that take into consideration the economic objectives at hand, despite of the wide use of Bayesian decision theory (Berger, 1985) and the increasing applications of Bayesian framework in finance, e.g., Kandel and Stambaugh (1996), Barberis (2000), Pastor and Stambaugh (2000), Brennan and Xia (2001), Avramov (2002), Cremers (2002), Cohen, Coval and Pástor (2005) and Wang (2005).

This paper explores a general approach to form objective-based priors in the context of an optimal portfolio selection problem. This problem has been extensively analyzed both theoretically and empirically ever since Markowitz's (1952) seminal mean-variance framework. Zellner and Chetty (1965), Klein and Bawa (1976) and Brown (1979) and Jorion (1986) are earlier studies that use the Bayesian approach to account for parameter uncertainty. Recently, Pastor (2000), Pastor and Stambaugh (2000), Avramov (2004) and Tu and Zhou (2004) study how the portfolio selection problem is impacted by an investor's varying prior beliefs on an asset pricing model and data-generating process. But the linkage between priors and the economic objective functions has not been addressed in neither these studies nor elsewhere in the finance literature.

The idea of this paper is to combine priors with the first-order condition of the portfolio maximization problem. Even in the absence of any data specific information, the investor's objective function places restrictions on the parameters of the model. As it turns out, such restrictions improve the Bayesian portfolio choice decision substantially.

The remainder of the paper is organized as follows. Section 2 provides the Bayesian framework. Section 3 extends the analysis to the case where the returns are predictable. Section 4 provides empirical results. Section 5 conducts a classical simulation comparison of the Bayesian procedures versus the classical ones. Section 6 concludes.

2 The Bayesian framework

1.1 The portfolio choice problem

Consider the standard portfolio choice problem where an investor chooses his optimal portfolio among N risky assets and a riskless asset. Let r_{ft} and r_t be the rates of returns on the riskless asset and N risky assets at time t, respectively. We define $R_t \equiv r_t - r_{ft}1_N$ as the excess returns, i.e., the returns in excess of the riskless asset, where 1_N is an N-vector of ones. The standard assumption on the probability distribution of R_t is that R_t is independent and identically distributed over time, and has a multivariate normal distribution with mean μ and covariance matrix V at any time t.

For simplicity, consider first the standard mean-variance framework, while the non-quadratic case will be discussed later. In the mean-variance framework, the investor is assumed to choose portfolio weights w so as to maximize the quadratic objective function

$$U(w) = E[R_{T+1}] - \frac{\tau}{2}\,\mathrm{Var}[R_{T+1}] = w'\mu - \frac{\tau}{2}w'Vw, \tag{1}$$

where τ is the coefficient of relative risk aversion. It is well-known that, when both μ and V are assumed known, the portfolio weights are

$$w^* = \frac{1}{\tau}V^{-1}\mu, \tag{2}$$

and the maximized expected utility is

$$U(w^*) = \frac{1}{2\tau}\mu'V^{-1}\mu = \frac{\theta^2}{2\tau}, \tag{3}$$

where $\theta^2 = \mu'V^{-1}\mu$ is the squared Sharpe ratio of the *ex ante* tangency portfolio of the risky assets.

However, w^* is not computable in practice because μ and V are unknown. To implement the above mean-variance theory of Markowitz (1952), the optimal portfolio weights are usually estimated by using a two-step procedure. First, the mean and covariance matrix of the asset returns are estimated based on the observed data. Second, these sample estimates are then treated as if they were the true parameters, and are simply plugged into (2) to compute the optimal portfolio weights. Kan and Zhou (2004) show that such plug-in approaches are not optimal and obtain better estimates for the optimal portfolio weights by incorporating the economic objective function into consideration. In this paper, we make use of the economic objective function by forming suitable priors in the Bayesian framework. In contrast with the classical statistics set-up, our Bayesian approach is more flexible because it allows easy adaption to different objective functions and to alternative data-generating processes.

1.2 Bayesian solution

Following Zellner and Chetty (1965), the Bayesian optimal portfolio is obtained by maximizing the expected utility under the predictive distribution, i.e.,

$$\begin{aligned} \hat{w}^{\text{Bayes}} &= \text{argmax}_w \int_{R_{T+1}} U(w) p(R_{T+1}|\mathbf{\Phi}_T)\, \mathrm{d}R_{T+1} \\ &= \text{argmax}_w \int_{R_{T+1}} \int_\mu \int_V U(w) p(R_{T+1}, \mu, V|\mathbf{\Phi}_T)\, \mathrm{d}\mu \mathrm{d}V \mathrm{d}R_{T+1}, \end{aligned} \tag{4}$$

where $U(w)$ is the utility of holding a portfolio w at time $T+1$, $p(R_{T+1}|\mathbf{\Phi}_T)$ is the predictive density, $\mathbf{\Phi}_T$ is the data available at time T, and

$$p(R_{T+1}, \mu, V|\mathbf{\Phi}_T) = p(R_{T+1}|\mu, V, \mathbf{\Phi}_T) p(\mu, V|\mathbf{\Phi}_T), \tag{5}$$

where $p(\mu, V|\mathbf{\Phi}_T)$ is the posterior density of μ and V. In contrast to the two-step solution that treats sample estimates as the true parameters, the Bayesian approach accounts for the estimation error. Brown (1976), Klein and Bawa (1976), and Stambaugh (1997), among others, choose the standard diffuse prior on μ and V,

$$p_0(\mu, V) \propto |V|^{-\frac{N+1}{2}}, \tag{6}$$

in the absence of using any data information. They show that the resulted optimal portfolio weights are generally slightly better than the classical plug-in approach, while Kan and Zhou (2005) verifies this analytically. Nevertheless, the differences between the two portfolio weights are known to be small and economically insignificant.

The closeness of the results between the classical and the Bayesian is not surprising from a statistical point of view. Neither the classical nor diffuse prior utilizes any information of data or the objective. Because of using the same diffuse information, their results should be close. On the other hand, as shown by Kan and Zhou (2005), the implied estimators of the Bayesian diffuse prior can be dominated by alternative estimators. This clearly says that the diffuse prior is not optimal to solving the portfolio problem. Then, the question is how to construct useful priors that can improve the investor's expected utility.

1.3 Priors incorporating objectives

Consider first the case in which no data information is available. In this case, we show that one can construct informative priors based on the objective function. Suppose that we are interested in forming a normal prior on μ,

$$\mu \sim N(\mu_0, \sigma_\mu^2 V_0), \tag{7}$$

where σ_μ is a scale prior parameter to indicate how tight the variance of μ is relative to V_0. To reflect the economic objective, it is natural to link the prior to the first-order condition (2),

$$\mu \sim N(V\tau w_0, \sigma_\mu^2 V_0) \tag{8}$$

where w_0 is our prior on the portfolio weights. This says that the prior mean is proportional to both the covariance matrix of the asset returns and the prior weights. Given the prior weights, the prior expected returns are high on those assets whose risks are high. Without using any data specific information, a simple prior on w_0 is an equal-weighted portfolio with weights summing to 0.5, which represents a strategy of investing 50% of the money into the risky assets.[1] This might be conservative to many. An alternative prior is to allow the weights summing to 1, which represents an aggressive strategy.

The above prior requires also the specification of V_0. A simple way of doing so is to use the identity matrix,

$$\mu \sim N(V\tau w_0, \sigma_\mu^2 I_N). \tag{9}$$

Note that σ_μ^2 reflects the degree of uncertainty about μ_0. A zero value of σ_μ^2 implies a dogmatic belief on μ_0 to be the true mean and there is no uncertainty. A value od $\sigma_\mu^2 = \infty$ suggests

[1]DeMiguely, Garlappiz, and Uppal (2005) provide a detailed discussion for this naive rule and find it is hard to beat.

that μ_0 is not informative at all about the true mean. Other than these two extremes, σ_μ^2 places modest informative prior beliefs on the degree of uncertainty of how μ_0 is close to be the true mean.

However, the identify matrix specification has an undesired property. It measures the difference of any value μ_d from μ_0,

$$\mu_d - \mu_0 \neq 0, \tag{10}$$

by placing equal importance on the deviations of each element of μ_d from μ_0. While this may be plausible in some applications, it does not measure adequately the investor's assessment of the deviations given his utility function. To see this, let w_d and w_0 be the portfolio weights associated with μ_d and μ_0 based on the objective function. It is easy to show that

$$U(w_0|\mu_0) - U(w_d|\mu_0) \approx \frac{1}{2}[\mu_d - \mu_0]'\Omega^{-1}[\mu_d - \mu_0]. \tag{11}$$

where

$$\Omega = \left\{ \left\{ \frac{\partial^2 U}{\partial w \partial \mu'}[w_0|\mu_0] \right\}' \left\{ \frac{\partial^2 U}{\partial w \partial w'}[w_0|\mu_0] \right\}^{-1} \left\{ \frac{\partial^2 U}{\partial w \partial \mu'}[w_0|\mu_0] \right\} \right\}^{-1}. \tag{12}$$

Hence, from the perspective of utility evaluation, the investor measures the importance of deviations with weights Ω^{-1}. This suggests a potentially better prior on μ is

$$\mu \sim N\left[V\tau w_0, \sigma_\mu^2 \left(\frac{1}{s^2}\Omega\right)\right], \tag{13}$$

where s^2 is the average of the diagonal elements of Ω. In this way, the investors's objective function, the utility function here, is linked to prior density. Note that this prior is invariant to any positive monotonic transformations of the utility function. In the case of mean-variance utility, it is easy to verify that $\Omega = V$. So, the above prior can be simply written as

$$\mu \sim N\left[V\tau w_0, \sigma_\mu^2 \left(\frac{1}{s^2}V\right)\right], \tag{14}$$

where, as we recall, V is the covariance matrix of the asset returns.

The objective-prior has an interesting relation to Black and Litterman's (1992) asset allocation method the received attention of many practitioners (see, e.g., Litterman (2003) and Meucci (2005)). They argue that, if w_0 is taken as the market value-weighted portfolio, $V\tau w_0$ is the equilibrium expected return as investors hold the market in equilibrium. It is this expected return that are used in their allocation decision which yields more balanced portfolios than the standard optimal mean-variance one. Our approach here allows it as a prior for Bayesian allocation that combines both the prior and the data. In contrast, Black and Litterman's Bayesian analysis is ad hoc and does not reply on the likelihood function of the data.

Consider now the case in which data is used to form informative priors on the parameters. For simplicity, assume we use 10 years of monthly data. Let $\hat{\mu}_{10}$ and $\hat{V}_{10}$ be the sample mean and covariance matrix, respectively. Then, the standard Bayesian informative prior on μ based on the 10 years data may be written as

$$\mu \sim N\left[\hat{\mu}_{10}, \sigma_\mu^2\left(\frac{1}{\hat{s}^2}\hat{V}_{10}\right)\right], \tag{15}$$

where $\hat{s}^2$ is the average diagonal elements of $\hat{V}_{10}$, and σ_μ^2 remain a scale parameter to indicate the degree of uncertainty. An alternative is the conjugate prior (Zellner, 1970),

$$\mu \sim N\left[\hat{\mu}_{10}, \sigma_\mu^2\left(\frac{1}{\hat{s}^2}\hat{V}\right)\right], \tag{16}$$

which replaces the earlier $\hat{V}_{10}$ by $\hat{V}$. In empirical empirical applications, the performance of these two priors are similar **Is this true?** and hence we consider only the first in what follows.

Given the data, an objective-based Bayesian starts from the non-data prior 14, updates his prior based on the 10 years data and then this for his future decision making. This is analogous to the updating the diffuse prior to get (15) and (16). The updated prior on μ is

$$\mu \sim N\left[\hat{\mu}_{10}^*, \sigma_\mu^2\left(\frac{1}{s^2}V\right)\right], \tag{17}$$

where $\hat{\mu}_{10}^* = V\tau\hat{w_{10}}$ and $\hat{w_{10}}$ is the Bayesian optimal portfolio weights based on the 10 years data. It is interesting that the conjugate prior provides a similar covariance structure to that of the objective-based prior. However, their means are entirely different. In applications, t the means are more difficult to estimate and it is them that make the most differences in portfolio decisions.

Pastor (2000) and Pastor and Stambaugh (2000) introduce a class of interesting priors that reflect investors' degree of belief in an asset pricing model. To see how this prior is formed, assume $R_t = (y_t, x_t)$, where y_t contains the excess returns of m non-benchmark positions and x_t contains the excess returns of K $(= N - m)$ benchmark positions. Consider a factor model multivariate regression,

$$y_t = \alpha + Bx_t + u_t, \tag{18}$$

where u_t is an $m \times 1$ vector of residuals with zero means and a non-singular covariance matrix $\Sigma = V_{11} - BV_{22}B'$, and α and B are related to μ and V, through

$$\alpha = \mu_1 - B\mu_2, \qquad B = V_{12}V_{22}^{-1}, \tag{19}$$

where μ_i and V_{ij} are the corresponding partition of μ and V,

$$\mu = \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \; V = \begin{pmatrix} V_{11} & V_{12} \\ V_{21} & V_{22} \end{pmatrix}. \tag{20}$$

For a factor-based pricing model, such as the three-factor model of Fama and French (1993), the restriction is $\alpha = 0$.

To allow for mispricing uncertainty, Pastor (2000) and Pastor and Stambaugh (2000) specify the prior distribution of α as a normal distribution conditional on Σ,

$$\alpha|\Sigma \sim N\left(0, \sigma_\alpha^2 \left(\frac{1}{s_\Sigma^2}\Sigma\right)\right), \tag{21}$$

where s_Σ^2 is a suitable prior estimate for the average diagonal elements of Σ. The above alpha-Sigma link is also explored by MacKinlay and Pástor (2000) in the classical framework. The magnitude of σ_α represents an investor's level of uncertainty about a given model's pricing ability. When $\sigma_\alpha = 0$, the investor believes dogmatically in the model and there is no mispricing uncertainty. On the other hand, when $\sigma_\alpha = \infty$, the investor believes that the pricing model is entirely useless.

In contrast with the priors motivated by asset pricing theory, the objective-based priors do not reply on the degree to which the asset pricing theory is true. What requires is only that the objective function of the price taking investor be known and well specified. The quadratic utility is assumed earlier, which can be relaxed as follow.

Consider how to obtain μ_0 and Ω in the power utility case which is one of the most popular utilities. The utility function is defined as

$$U[w|\mu[= E\left[\frac{[1 + R_f + wR_{T+1}]^{1-\gamma}}{1-\gamma}|\mu\right]. \tag{22}$$

The first-order condition is

$$E\left[[1+R_f+wR_t]^{-\gamma}R_t|\mu\right]=0, \tag{23}$$

from which we can solve for μ for any given w_0. As there is no analytical solutions, we can either use the earlier normal distribution of μ with mean μ_0 as an approximation or truncated the posterior draws of μ in a neighborhood of equation (refeq:power-foc). The second-order derivatives are:

$$\frac{\partial U}{\partial w}[w|\mu]=E\left[[1+R_f+wR_t]^{-\gamma}R_t|\mu\right], \tag{24}$$

and

$$\frac{\partial^2 U}{\partial w\partial w'}[w|\mu]=-\gamma E\left[[1+R_f+wR_t]^{-\gamma-1}R_tR_t'|\mu\right]. \tag{25}$$

Although Ω now is also not available analytically, it can be computed numerically. Samples of μ nay be drawn, and then we can approximate the second-order derivatives by

$$\frac{\partial^2 U}{\partial w\partial \mu'}[w_0|\mu_0]=\left[\frac{\partial f(\mu_0)}{\partial \mu_1},\frac{\partial f(\mu_0)}{\partial \mu_2},\cdots,\frac{\partial f(\mu_0)}{\partial \mu_N}\right], \tag{26}$$

where

$$\frac{\partial f(\mu_0)}{\partial \mu_i}=\left[f(\mu_0-2he_i)-8f(\mu_0-he_i)+8f(\mu_0+he_i)-f(\mu_0+2he_i)\right]/(12h), \tag{27}$$

where $f(\mu_0)\equiv\frac{\partial U}{\partial w}[w_0|\mu_0]$, $e_i=(0,0,\cdots,0,1,0,\cdots,0,0)$ is an $N\times 1$ vector with 1 as the ith element and zero otherwise, and h is a small number, such as 10^{-6}.

3 Priors when returns are predictable

Kandel and Stambaugh (1996) and Barberis (2000) show that incorporating return predictability plays an important role in portfolio decisions. Avramov (2004) extends this in a multivariate asset setting. The question we address here is whether objective-based prior can still result significant economic gains in the presence of predictability.

Following earlier studies, we assume that the returns are related to L predictive variables by linear regression,

$$R_t=\mu+\mu_1 Z_{t-1}+v_t, \tag{28}$$

where $v_t \sim N(0, \Sigma_{rr})$ and the predictive variables Z_{t-1} have an VAR process,

$$Z_t = \psi + \psi_1 Z_{t-1} + u_t, \tag{29}$$

with $u_t \sim N(0, \Sigma_{zz})$. Let $y_t' = [R_t', Z_{t-1}']$, $x_t = (1, y_{t-1}')'$, and Y and X be $(N + L) \times T$ and $(N + L + 1) \times T$ matrices of the y_t's and x_t's, respectively, then the model can be written in the standard form,

$$Y = XB + U, \tag{30}$$

where U are formed by the corresponding residuals, $\text{vec}(U) \sim N(0, \Sigma \otimes I_T)$, where

$$\Sigma = \begin{pmatrix} \Sigma_{rr} & \Sigma_{rz} \\ \Sigma_{zr} & \Sigma_{zz} \end{pmatrix}. \tag{31}$$

We in what follows focuses on how to impose an objective-based prior since the solution to the Bayesian optimal portfolio problem is known in the diffusion prior case. Conditional on any given prior for μ_1, the mean and variance conditional on Z_{T-1} are relevant for the usual utility maximization. Similar to the iid case, we obtain an objective-based prior on μ as

$$\mu \sim N\left[\Sigma_{rr}\tau w_0, \sigma_\mu^2\left(\frac{1}{s^2}\Sigma_{rr}\right)\right]. \tag{32}$$

This prior can easily be combined with a prior on predictability and on "no-predictability" as set forth in Kandel and Stambaugh (1996).

4 Empirical results

In this section, we compare the objective-based priors with its usual alternatives. The data are monthly returns of the well-known Fama-French 25 book-to-market and size portfolios and their three factors from January 1965 to December 2004.[2]

Consider first the possible gains of switching from diffuse priors to the objective-based prior without using any data information. Following Kandel and Stambaugh (1996) and Pástor and Stambaugh (2000), the utility gain of switching from the diffuse prior to the objective-based prior is the difference of expected utilities under the posterior distribution of

[2]We are grateful to Ken French for making this data available on his website.

the latter. Let E^* and V^* be the predictive mean and variance-covariance matrix of the asset returns under objective-based prior without data, equation (14), and w_O be the associated optimal portfolio allocation. Then the expected utility is given by

$$EU_O = w_O^{'} E^* - \frac{1}{2} A w_O^{'} V^* w_O, \tag{33}$$

where A is the degree of relative risk aversion. The allocation, w_D, which is optimal under the diffuse prior, should have an expected utility of

$$EU_D = w_D^{'} E^* - \frac{1}{2} A w_D^{'} V^* w_D. \tag{34}$$

Notice that this expected utility is evaluated based on the same E^* and V^* of the objective-based prior. Because of this, the difference

$$CE = EU_O - EU_D \tag{35}$$

is interpreted as the 'perceived' loss in terms of certainty-equivalent return to an investor who is forced to accept the optimal portfolio selection based on the diffuse prior. Since w_O is optimal under the objective-based prior, CE is always positive or zero by construction. The issue is how big this value can be. Generally speaking, values over a couple of percentage points per year are deemed as economically significant.[3]

Table 1 reports the results. When we apply the objective-based prior to 10 years worth of data ($T = 60$), the utility gains are overwhelmingly large. They range from an annual rate of 22% to 124%. The large difference is driven by the fact that the posterior returns are very sensitive to prior specifications when $T = 60$. This can also understood by the simulation results of Kan and Zhou (2005) who show that, with a sample size of $T = 60$, the estimated parameters can be far away from the true ones. Even with a sample size of $T = 480$, there is still substantial variability in the estimates. The objective-based prior seems to provide an informative way to shift the posterior mean from the sample one to a more reasonable value.

As sample size grows, the influence of priors decreases because the posterior distribution is completely determined by the data when the sample size is infinity. However, with a

[3]Fleming, Kirby and Ostdiek (2001) provides a similar but different measure in the classical framework.

sample size of as large as $T = 480$, Table 1 shows that the utility gains are still substantial. At $\sigma_\mu = 1\%$, the gain is greater than 8%, although it eventually decreases to an insignificant amount of 0.04% at $\sigma_\mu = 5\%$. Overall, it is clear that the objective-based prior make a significant economic difference in portfolio selections.

Consider now the case in which some of the data is used to form informative priors. In this case, the data prior, equation (15) plays the role of earlier diffuse prior, while the comparable objective-based prior is the one given by equation (17) which updates the previous one with the same length of data. Table 2 **Since we have in total 480 data, should T goes only to 360 in table 2?** provides the results. The objective-based prior outperforms the data-based prior substantially when $T < 240$, or when $\sigma_\mu \geq 2\%$. However, the gains are less dramatic than the diffuse prior case, and they gains are not necessarily smaller or greater as T increases. For example, quite a few of the gains when $T = 480$ are even greater than those with few samples. There are two explanations for this. First, in a given application, the first 10 years of data may not necessarily informative to the samples that follow, and hence the expected utilities may not be a monotonic function of the sample size. Second, even if they are, their difference may not necessarily be so.

Consider finally the performance of the objective-based priors relative to those based on asset pricing models. Taking x_t as the Fama-French three-factors, the degree of belief on the validity of the Fama-French three-factor model is represented by the alpha prior, equation (21). For simplicity, we assume $\sigma_\alpha^2 = \sigma_\mu^2$ in the comparison. Table 3 provides the results. Similar to the data-based prior case, the utility gains are economically significant for all the sample sizes when $\sigma_\mu \geq 2\%$.

In summary, maximizing a utility function provide a useful guidance for choosing priors in Bayesian decision making. Using the Fama-French data, we find that such objective-based priors outperform both standard statistical priors and asset-pricing-based ones significantly. Even with sample size as large as $T = 480$, the utility gains are still economically significant.

5 A comparison in the classical framework

Because the true parameters of the data-generating process is unknown in the real world, the Bayesian utility comparison in the previous section is based on E^* and V^*, the mean and variance estimated based on the objective-based prior with or without data information. The gains are positive by design. One might argue that this benchmark is unfairly biased against the diffuse prior, and the diffuse prior might perform better under the true parameters.

To answer this question, we conduct the following simulation experiment. First, we set true parameters of the model as the sample mean and covariance matrix of the monthly returns of the Fama-French data from January 1965 to December 2004. Then, we simulate 1000 data sets from the assumed normal distribution of asset returns, and compute the expected utilities based on the diffuse prior and the objective-based priors. In contrast with earlier Bayesian analysis, we evaluate the expected utilities at the assumed true parameters so that no bias toward neither priors exists. Table 4 reports the average utility gains over the simulated data sets. It is striking that the objective-based prior achieve even greater gains over the diffuse prior when $T = 60$, suggesting that the diffuse prior is indeed inadequate with small sample size. However, as sample size increases, the magnitude of the gains, though economically significant, decrease substantially. Nevertheless, even when sample size T is as large as $T = 480$, the gains are still economically important.

One may view Bayesian portfolio decisions, based on either the diffuse prior or objective-based priors, as purely decision rules that are functions of the samples. In the classical framework, their performance is judges by the expected values of the functions or by the simulated sample averages. By this classical statistics criterion, the objective-based Bayesian analysis is clearly a better decision tool than the diffuse one. This is also true, as shown by Table 4 and 5, when we compare the objective-based prior with the data- and asset-pricing-based ones.

6 Conclusion

This paper explores the link between Bayesian priors and economic objective functions. Once incorporating the economic objectives into priors, we find that the gains are substantial, and sometimes are enormous. The superior performance is present not only in using the real data, but also in simulations.

Although our study focuses on portfolio choice, the results do suggest economic objective-based priors should be explored in almost any financial decision making, especially where the Bayesian framework is deemed appropriate because such priors contains information on both the nature of the data and the plausible parameters of the model.

References

Avramov, D., 2002, Stock return predictability and model uncertainty. Journal of Financial Economics 64, 423–458.

Avramov, D., 2004, Stock predictability and asset pricing models, *Review of Financial Studies* 17, 699–738.

Barberis, N., 2000, Investing for the long run when returns are predictable. Journal of Finance 55, 225–264.

Bawa, Vijay S., Stephen J. Brown, and Roger W. Klein, 1979, *Estimation Risk and Optimal Portfolio Choice* (Amsterdam: North-Holland).

Berger, James O., 1985, *Statistical Decision Theory and Bayesian Analysis* (Springer-Verlag, New York).

Black, F., Litterman, R., 1992, Global portfolio optimization, Financial Analysts Journal 48, 28–43.

Blattberg, Robert C., and Nicholas J. Gonedes, 1974, A comparison of the stable and Student distributions as statistical models for stock prices, *Journal of Business* 47, 244–280.

Brown, Stephen J., 1976, Optimal portfolio choice under uncertainty, Ph.D. dissertation, University of Chicago.

Brown, Stephen J., 1978, The portfolio choice problem: comparison of certainty equivalence and optimal Bayes portfolios, *Communications in Statistics — Simulation and Computation* 7, 321–334.

Cohen, Randolph B., Joshua D. Coval, Lubos Pástor, 2005, Judging fund managers by the company they keep, *Journal of Finance* 60, 1057–1096.

Cremers, M., 2002, Stock return predictability: A Bayesian model selection perspective. Review of Financial Studies 15, 1223–1249.

DeMiguely, Victor, Lorenzo Garlappiz, and Raman Uppal, 2005, How Inefficient is the $1/N$ Asset-Allocation Strategy? Working paper, London Business School and The University of Texas at Austin.

Fama E.F., French, K.R., 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Fleming, J., Kirby, C., Ostdiek, B., 2001, The economic value of volatility timing, *Journal of Finance* 56, 329–352.

Jorion, Philippe, 1986, Bayes-Stein estimation for portfolio analysis, *Journal of Financial and Quantitative Analysis* 21, 279–292.

Kan, R., Zhou, G., 2005, Optimal Estimation for Economic Gains: Portfolio Choice with Parameter Uncertainty, working paper.

Kandel, S., Stambaugh, R.F., 1996, On the predictability of stock returns: An asset-allocation perspective, *Journal of Finance* 51, 385–424.

Klein, Roger W., and Vijay S. Bawa, 1976, The effect of estimation risk on optimal portfolio choice, *Journal of Financial Economics* 3, 215–231.

Lehmann, E.L., and George Casella, 1998, *Theory of Point Estimation* (Springer-Verlag, New York).

Litterman, B., 2003, Modern Investment Management: An Equilibrium Approach, Wiley, New York.

Markowitz, Harry M., 1952, Mean-variance analysis in portfolio choice and capital markets, *Journal of Finance* 7, 77–91.

Meucci, A., 2005, Risk and Asset Allocation, Springer-Verlag, New York.

Pástor, Ľ., 2000, Portfolio selection and asset pricing models, *Journal of Finance* 55, 179–223.

Pástor, Ľ., Stambaugh, R.F., 2000, Comparing asset pricing models: an investment perspective, *Journal of Financial Economics* 56, 335–381

Tu, J., and G. Zhou, 2004, Data-generating process uncertainty: what difference does it make in portfolio decisions? Journal of Financial Economics 72, 385–421.

Wang, Zhenyu, 2005, A shrinkage approach to model uncertainty and asset allocation, *Review of Financial Studies* 18, 673–705.

Zellner, A., Chetty, V.K., 1965, Prediction and decision problems in regression models from the Bayesian point of view, *Journal of the American Statistical Association* 60, 608–616.

Table 1: Utility gains of switching from diffuse to objective-based priors

The table reports the utility gains (annualized) of switching from the standard diffuse prior,

$$p_0(\mu, V) \propto |V|^{-\frac{N+1}{2}}$$

to the objective-based prior

$$p_0(\mu, V) \propto N\left[\tau V w_0, \sigma_\mu^2 \left(\frac{1}{s^2} V\right)\right] \times |V|^{-\frac{N+1}{2}},$$

where σ_μ^2 reflects the degree of uncertainty about μ_0 and w_0 is proportional to a constant with $\sum w_{0i} = 0.5$ or 1, respectively.

$\sum w_{0i}$	T	σ_μ				
		1%	2%	3%	4%	5%
			T= 60			
0.5	60	123.93	90.23	58.47	36.17	22.36
1	60	125.47	91.36	59.18	36.68	22.66
			T= 120			
0.5		76.97	31.80	12.54	5.42	2.61
1		75.57	31.20	12.33	5.33	2.57
			T= 180			
0.5		53.45	14.62	4.54	1.74	0.78
1		52.54	14.39	4.45	1.71	0.77
			T= 240			
0.5		38.72	8.45	2.36	0.85	0.38
1		38.00	8.27	2.33	0.84	0.37
			T= 360			
0.5		15.95	2.66	0.67	0.24	0.10
1		15.28	2.54	0.65	0.22	0.10
			T= 480			
0.5		8.87	1.26	0.31	0.10	0.04
1		8.70	1.24	0.30	0.10	0.04

Table 2: Utility gains of switching from data-based to objective-based priors

The table reports the utility gains (annualized) of switching from the data-based prior

$$p_0(\mu, V) \propto N\left[\hat{\mu}_{10}, \sigma_\mu^2\left(\frac{1}{\hat{s}^2}\hat{V}_{10}\right)\right] \times |V|^{-\frac{\nu_V+N+1}{2}} exp\left\{-\frac{1}{2}trHV^{-1}\right\}$$

to the objective-based prior

$$p_0(\mu, V) \propto N\left[\hat{\mu}^*_{10}, \sigma_\mu^2\left(\frac{1}{s^2}V\right)\right] \times |V|^{-\frac{\nu_V+N+1}{2}} exp\left\{-\frac{1}{2}trHV^{-1}\right\},$$

where $\hat{\mu}^*_{10} = V\tau\hat{w_{10}}$, $\hat{w_{10}}$ is the Bayesian optimal portfolio weights based on the 10 years data when w_0 is proportional to a constant with $\sum w_{0i} = 0.5$ or 1, respectively, σ_μ^2 reflects the degree of uncertainty about μ_0, $H = T_{10}\hat{V}_{10}$, $\nu_V = T_{10}$, $T_{10} = 120$ (sample size of 10 year monthly returns), and $\hat{V}_{10}$ is the sample covariance of the previous 10 year monthly returns.

Sum of Prior weights	σ_μ				
	1%	2%	3%	4%	5%
			T= 60		
0.5	54.49	30.32	18.26	12.32	8.88
1	53.23	29.46	17.94	12.31	8.87
			T= 120		
0.5	44.84	44.96	31.38	19.28	12.31
1	42.61	44.07	31.31	19.12	12.22
			T= 180		
0.5	34.40	14.34	7.44	4.07	2.34
1	33.70	14.31	7.40	4.23	2.36
			T= 240		
0.5	17.84	4.30	1.62	0.73	0.39
1	17.27	4.32	1.66	0.70	0.36
			T= 360		
0.5	6.61	1.97	0.78	0.35	0.18
1	6.36	1.90	0.75	0.34	0.18
			T= 480		
0.5	42.91	8.71	2.36	0.95	0.47
1	43.05	8.73	2.36	0.94	0.47

Table 3: Utility gains of switching from Fama-French three-factor model-based to objective-based priors

The table reports the utility gains (annualized) of switching from a prior reflecting the degree of belief in the Fama-French three-factor model,

$$p_0(\alpha, B, V) \propto N(B\mu_2, \sigma_\mu^2 \frac{1}{s_\Sigma^2}\Sigma) \times |V|^{-\frac{N+1}{2}},$$

where $\Sigma = V_{11} - V_{12}V_{22}^{-1}V_{21}$ and s_Σ^2 is the average of the diagonal elements of Σ, to the objective-based prior

$$p_0(\mu, V) \propto N\left[\tau V w_0, \sigma_\mu^2\left(\frac{1}{s^2}V\right)\right] \times |V|^{-\frac{N+1}{2}},$$

where σ_μ^2 reflects the degree of uncertainty about μ_0 and w_0 is proportional to a constant with $\sum w_{0i} = 0.5$ or 1, respectively.

Sum of Prior weights	σ_μ				
	1%	2%	3%	4%	5%
			T= 60		
0.5	84.32	120.18	125.20	114.53	102.71
1	83.19	118.94	123.68	113.31	101.47
			T= 120		
0.5	42.53	39.94	26.21	18.26	14.11
1	40.26	38.64	25.49	17.92	13.82
			T= 180		
0.5	33.86	18.97	10.05	6.53	4.93
1	32.65	18.51	9.83	6.41	4.85
			T= 240		
0.5	28.62	11.63	5.54	3.45	2.57
1	27.68	11.30	5.43	3.39	2.52
			T= 360		
0.5	14.76	4.04	1.72	1.04	0.76
1	13.92	3.83	1.65	0.99	0.74
			T= 480		
0.5	8.42	1.87	0.74	0.43	0.30
1	8.16	1.81	0.72	0.41	0.30

Table 4: Monte Carlo utility gains of switching from diffuse to objective-based priors

This table reports the average utility gains of switching from diffuse prior to objective-based priors with data sets simulated from a multivariate normal distribution whose 'true' mean and covariance matrix, E^* and V^*, are calibrated from the monthly returns of the Fama-French 25 assets and the associated three factors from January 1965 to December 2004. The number of simulated data sets is 1000.

Prior types	σ_μ				
	1%	2%	3%	4%	5%
			T= 60		
0.5	185.14	185.19	168.42	143.16	118.04
1	186.06	185.77	168.78	143.39	118.11
			T= 120		
0.5	42.50	44.95	35.16	25.68	18.87
1	43.21	45.25	35.28	25.77	18.89
			T= 180		
0.5	18.99	21.88	15.67	10.80	7.66
1	19.55	22.07	15.79	10.84	7.65
			T= 240		
0.5	10.08	13.03	8.91	5.93	4.11
1	10.54	13.16	8.92	5.97	4.13
			T= 360		
0.5	3.64	6.18	3.97	2.56	1.76
1	3.97	6.25	3.99	2.57	1.75
			T= 480		
0.5	1.33	3.52	2.19	1.39	0.95
1	1.56	3.55	2.20	1.39	0.96

Table 5: Monte Carlo utility gains of switching from data-based to objective-based priors

This table reports the average utility gains of switching from the data-based to objective-based priors with data sets simulated from a multivariate normal distribution whose 'true' mean and covariance matrix, E^* and V^*, are calibrated from the monthly returns of the Fama-French 25 assets and the associated three factors from January 1965 to December 2004. The number of simulated data sets is 100.

Sum of Prior weights			σ_μ		
	1%	2%	3%	4%	5%
			T= 60		
0.5	71.52	98.15	87.66	68.40	52.92
1	72.32	98.58	87.56	68.28	52.60
			T= 120		
0.5	21.38	20.07	13.53	9.02	6.16
1	21.97	20.24	13.67	8.96	6.44
			T= 180		
0.5	16.38	9.49	5.15	2.94	2.00
1	16.61	9.41	5.10	3.04	2.04
			T= 240		
0.5	12.77	5.16	2.38	1.42	0.89
1	12.97	5.21	2.38	1.34	0.85
			T= 360		
0.5	8.04	1.81	0.67	0.40	0.14
1	8.11	1.80	0.69	0.41	0.23
			T= 480		
0.5	4.70	0.70	0.26	0.16	0.14
1	4.73	0.67	0.28	0.16	0.12

Table 6: Monte Carlo utility gains of switching from Fama-French three-factor model-based to objective-based priors

This table reports the average utility gains of switching from the Fama-French three-factor model-based to objective-based priors with data sets simulated from a multivariate normal distribution whose 'true' mean and covariance matrix, E^* and V^*, are calibrated from the monthly returns of the Fama-French 25 assets and the associated three factors from January 1965 to December 2004. The number of simulated data sets is 1000.

Sum of Prior weights	σ_μ				
	1%	2%	3%	4%	5%
			T= 60		
0.5	53.47	187.62	237.01	242.54	233.47
1	54.39	188.21	237.37	242.78	233.54
			T= 120		
0.5	21.33	55.66	55.95	50.49	45.63
1	22.04	55.96	56.07	50.58	45.66
			T= 180		
0.5	9.67	25.95	23.39	19.92	17.45
1	10.23	26.14	23.51	19.96	17.44
			T= 240		
0.5	5.22	15.58	13.33	11.06	9.58
1	5.68	15.71	13.34	11.10	9.60
			T= 360		
0.5	1.15	6.87	5.41	4.27	3.61
1	1.48	6.94	5.43	4.29	3.60
			T= 480		
0.5	0.26	4.24	3.32	2.67	2.30
1	0.49	4.28	3.33	2.67	2.30