Stock Returns and Inflation Risk: A Bayesian Approach

Tomek Katzur* Laura Spierdijk†

March 7, 2011

*Corresponding author. University of Groningen, Faculty of Economics & Business, Department of Economics, Econometrics & Finance, P.O. Box 800, NL-9700 AV Groningen, The Netherlands. Email: t.m.katzur@rug.nl.

†Laura Spierdijk, University of Groningen and Netspar. Address: Faculty of Economics & Business, Department of Economics, Econometrics & Finance, P.O. Box 800, NL-9700 AV Groningen, The Netherlands. Email: l.spierdijk@rug.nl. The authors would like to thank Laurens Swinkels, Marno Verbeek and Giovanna Nicodano for useful comments. We are also grateful to the participants of the Inquire Europe Seminar in Berlin, the XIX Conference at the University of Rome Tor Vergata and the Netspar International Pension Workshop in Amsterdam. The usual disclaimer applies.

Abstract

A widely adopted view in the economic literature is that an asset is a good hedge against inflation if the Fisher hypothesis holds true; i.e. if the marginal effect of a unit change in inflation on nominal returns (often referred to as the Fisher coefficient) is equal to unity. In a regression of volatile stock returns on slowly moving inflation rates the estimated Fisher coefficient is often characterized by large standard errors, due to which the Fisher hypothesis cannot be rejected. Investors who base their optimal investment portfolios on the assumption that stocks are a good hedge against inflation ignore the information contained in the Fisher coefficient. This raises the question whether stocks are still viewed as a good hedge against inflation if we take into account the parameter uncertainty involved with the estimated Fisher coefficient. This paper seeks an answer to this question by adopting a Bayesian approach. In our empirical study we find little traditional statistical evidence against the Fisher hypothesis, suggesting that stocks are a good hedge against both expected and unexpected inflation. By contrast, the Bayesian approach reveals a substantial exposure of stock returns to inflation risk.

Keywords: Fisher hypothesis, parameter uncertainty

JEL classification: G11, G14

1 Introduction

Inflation risk is one of the primary concerns for long-term investors such as pension funds. Although the short-run effects of an increase in the price level may seem small and negligible, the long-run effects of inflation can be substantial. With a relatively low annual inflation rate of 2%, a value of 100 dollar today is worth only 67 dollar in twenty years time. Inflation-linked instruments provide a hedge against inflation, but their availability is limited. Moreover, real returns on such assets are usually low. It may therefore be attractive for investors to extend their portfolio with an investment in stocks in order to benefit from the equity premium. However, stocks are potentially exposed to inflation risk.

A widely adopted view in the economic literature is that an asset is a good hedge against inflation if the 'Fisher hypothesis' holds true; see e.g. Fama and Schwert (1977), Boudoukh and Richardson (1993), Barnes et al. (1999), and Bekaert and Wang (2010). In his classical *Theory of Interest* (1930), Irving Fisher postulated that the anticipated rate of inflation is completely incorporated in the ex ante nominal interest rate. At the same time, he precluded a relation between the expected real rate and expected inflation, emphasizing the independence of the real and monetary sectors. The proposition that ex ante nominal returns contain the market's perception of anticipated inflation rates can be applied to all assets. As a consequence, expected nominal returns on any asset would move one-to-one with expected inflation. The marginal effect of a unit change in inflation on nominal returns is often referred to as the 'Fisher coefficient'. The Fisher hypothesis is alternatively formulated as real asset returns being statistically uncorrelated with expected inflation.

Empirical studies based on monetary assets produced ambiguous results about the Fisher effect (Roll, 1972). However, using the argument that stocks are claims to real assets, the Fisher effect was widely believed to hold for common stocks until the early seventies (Lintner, 1973; Fama and MacBeth, 1974; Nelson, 1976). This 'accepted dogma' (Fama and Macbeth, 1974) was subjected to serious empirical scrutiny only after the subsequent episode of soaring inflation rates and poor stock market performance (Jaffe and Mandelker, 1976; Nelson, 1976). Fama and Schwert (1977) translated the Fisher hypothesis into a regression framework and estimated the relation between stock returns and proxies of expected and unexpected inflation. Contrary to other assets, such as real estate, stock returns were found to be a poor hedge against both expected and unexpected inflation for the 1953 – 1971 period in the United States. These results were confirmed for other major stock markets by e.g. Solnik (1983) and Gultekin (1983). Instead of being an inflation hedge, stock holdings turned out to suffer from considerable exposure to inflation risk. Boudoukh and Richardson (1993) partially rehabilitate the Fisher hypothesis, however. They find evidence in favor of the Fisher effect for five-year stock returns. With a one-year investment horizon there is some evidence for a significantly negative relation between nominal stock returns and inflation. Solnik and Solnik (1997) extend the analysis of Boudoukh and Richardson (1993) to countries other than the United States. Using a panel data set consisting of stock index returns and inflation rates in eight major economies, they establish a Fisher coefficient that increases towards unity as the investment horizon gets longer.

Stock returns and inflation rates are usually characterized by different time series properties, the former being much more volatile than the latter in the short run (Bodie, 1976 and Schotman and Schweizer, 2000). Consequently, estimated Fisher coefficients often feature huge standard errors, particularly in short samples, due to which the Fisher hypothesis cannot be rejected. Investors may therefore believe that stocks are a good hedge against inflation, with important consequences for portfolio choice. It has been demonstrated in the asset allocation literature that a statistically significant relation in regressions of stock returns on predictor variables is not a necessary condition for economically significant results (see e.g. Kandel and Stambaugh, 1996). But investors who base their optimal investment portfolios on the assumption that stocks are a good hedge against inflation ignore the information contained in the Fisher coefficient. This raises the question whether stocks are still viewed as a good hedge against inflation if we take into account the parameter uncertainty involved with the estimated Fisher coefficient. This paper seeks an answer to this question by adopting a Bayesian approach.

The starting point of this study is a stylized investment problem, consisting of a long-term investor who divides his wealth between stocks and inflation-linked bonds paying a risk-free real rate. He sets his portfolio weights in order to maximize the expected utility associated with his real wealth at the end of his investment horizon. An important assumption the investor has to make is about the relation between real stock returns and inflation. He can a priori assume that real returns are independent of expected or unexpected inflation (or both), but he can also remain more agnostic by allowing for feedback between real returns and inflation. The investor who assumes that real stock returns are independent of inflation lives in a world free of inflation risk (where the Fisher hypothesis holds true), whereas the agnostic investor is exposed to inflation risk via his stock holdings. In recent years the stock return predictability literature has been using methods that are capable of dealing with the problem of parameter uncertainty in investment decisions; see for example Barberis (2000). These studies adopt a Bayesian approach to calculate optimal asset allocations in the presence of parameter uncertainty. To assess if stocks are a good hedge against inflation we apply this Bayesian approach. We analyze the difference in stock allocations between the agnostic investor and the investor who makes strong a priori assumptions about the relation between real stocks returns and inflation. In a world free of inflation risk the benchmark and agnostic investors would have the same optimal portfolio weights. Hence, the difference in portfolio weights between the two investors provides a new measure of the inflation risk exposure of stocks that accounts for parameter uncertainty.

The setup of the remainder of this paper is as follows. The stylized investment problem is described in Section 2. Section 3 explains the Bayesian approach to determine optimal portfolio weights in the presence of parameter uncertainty. The data used for the empirical part of this paper are described in Section 4. Section 5 contains the empirical results. Section 6 discusses several robustness checks, among which an extension of our analysis that takes into account the dividend yield. Finally, Section 7 concludes.

2 Theoretical framework

Our starting point is the pure asset allocation problem studied by, amongst others, Barberis (2000), Ang and Bekaert (2002) and Guidolin and Timmerman (2007). We consider two investors who both have initial nominal wealth $\widetilde{W}_t = 1$ at time t, when the price level equals $P_t = 1$. Each investor seeks to maximize utility over real-term wealth $W_{t+k} = \widetilde{W}_{t+k}/P_{t+k}$ at time $t + k$

(where k denotes e.g. months). We assume power utility over real term wealth; that is

$$u(W_{t+k}) = \frac{W_{t+k}^{1-\phi}}{1-\phi}, \tag{1}$$

where ϕ is the coefficient of relative risk aversion. Power utility (also referred to as constant relative risk aversion utility) is a widely applied utility function, with encompasses both quadratic and logarithmic utility functions as special cases (Wakker, 2008). At time t the investor decides about the proportion of wealth λ_t he wishes to allocate to a stock index, the other investment option being a riskless inflation-linked bond with maturity k. He is assumed to hold these investments until time $t + k$ ('buy and hold'). Although the inflation-linked bond provides a hedge against inflation, its real return is usually low. To benefit from the real equity premium, it may be attractive for the investor to extend his portfolio with an investment in a stock index. The utility of terminal real wealth W_{t+k} can be written in terms of real stock returns as

$$u(W_{t+k}) = \frac{\left[\lambda \exp(r_t(k)) + (1-\lambda)\exp(r_{f,t}(k))\right]^{1-\phi}}{1-\phi} \tag{2}$$

where $r_{f,t}(k)$ is the continuously compounded k-period risk-free real rate and $r_t(k) = \sum_{s=1}^{k} r_{t+s}$ the k-period logarithmic real stock return, which boils down to the sum of one-period log real returns.

It is our goal to determine the impact of inflation risk on the investor's optimal choice. In the seminal work of Fama and Schwert (1977) the relation between asset returns and inflation has been studied using linear regressions of the form

$$R_{t+1} = \mu + \beta \mathbb{E}_t[\pi_{t+1}] + \gamma\big(\pi_{t+1} - \mathbb{E}_t[\pi_{t+1}]\big) + \varepsilon_{t+1}, \tag{3}$$

where R_{t+1} denotes nominal asset returns from time t to $t + 1$, $\mathbb{E}_t[\pi_{t+1}]$ one-period expected inflation, and $\pi_{t+1} - \mathbb{E}_t[\pi_{t+1}]$ one-period unexpected inflation. An asset is called a complete hedge against inflation if $\beta = \gamma = 1$. In this case real returns are uncorrelated with inflation and nominal asset returns move one-to-one with total inflation. This situation corresponds to Fisher's idea that the price of a stock, which ultimately represents a claim on real assets, should not be affected by inflation. Another possibility that has received considerable attention in the literature is that stock returns are affected by unexpected inflation only, that is $\beta = 1$ but $\gamma \neq 1$. Fama and Schwert (1977) call such an asset a complete hedge against expected inflation. This application of the Fisher (1930) hypothesis to stock returns has been studied empirically by, amongst others, Boudoukh and Richardson (1993) and Solnik and Solnik (1997).

Relating to this classical framework, we assume that our investor uses the following stylized reduced-form vector autoregressive (VAR) model to capture the dynamics between one-period real stock returns (r_{t+1}) and one-period expected (π^e_{t+1}) inflation and unexpected inflation (π^u_{t+1}):

$$\begin{pmatrix} r_{t+1} \\ \pi^e_{t+1} \\ \pi^u_{t+1} \end{pmatrix} = \begin{pmatrix} \mu_1 \\ \mu_2 \\ 0 \end{pmatrix} + \begin{pmatrix} 0 & \beta_1 & 0 \\ 0 & \beta_2 & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} r_t \\ \pi^e_t \\ \pi^u_t \end{pmatrix} + \begin{pmatrix} \varepsilon_{1,t+1} \\ \varepsilon_{2,t+1} \\ \varepsilon_{3,t+1} \end{pmatrix}, \tag{4}$$

where $(\varepsilon_{1,t}, \varepsilon_{2,t}, \varepsilon_{3,t})$ is a series of independent multivariate normally distributed disturbances, with mean zero and covariance matrix $\Sigma = \mathbb{C}\text{ov}\,(\varepsilon_{1,t}, \varepsilon_{2,t}, \varepsilon_{3,t})$. Here Σ has diagonal elements σ_i^2 for $i = 1, 2, 3$ and off-diagonal elements $\sigma_{ij}^2 = \sigma_{ji}^2$ for $i \neq j$. The variables $\pi^e_{t+1} = \mathbb{E}_t[\pi_{t+1}]$ and $\pi^u_{t+1} = \pi_{t+1} - \mathbb{E}_t[\pi_{t+1}]$ in Equation (4) are (proxies of) expected and unexpected inflation, respectively. The first equation in Model (4) directly relates real stock returns to expected inflation via the Fisher coefficient β_1. The correlation between $\varepsilon_{1,t}$ and $\varepsilon_{1,3}$ (say ρ_{13}) captures the influence of unexpected inflation on innovations in stock returns. Additionally, the VAR model allows for correlation between innovations in real stock returns and shocks in expected inflation via ρ_{12}. The second equation specifies expected inflation as an AR(1)-process. Finally, in the third equation unexpected inflation is assumed to be a white noise process with variance σ_3^2. The main difference between Model (4) and existing models in the literature (see e.g. Schotman and Schweitzer, 2000) is that we assume that expected inflation, rather than inflation itself, follows an AR(1)-process. This specification is adopted mainly for the sake of obtaining a more accurate model of expected inflation, as suggested by Ang et al. (2007).

2.1 Investor beliefs

In our subsequent analysis we quantify the impact of inflation risk on asset allocation by solving the investor's optimization problem for three different sets of beliefs about the relation between real stock returns and inflation in the model of Equation (4). These beliefs correspond to different definitions of inflation hedge, as proposed by Fama and Schwert (1977). First, we solve the problem for a benchmark investor who believes that stock returns act as a complete hedge against inflation, meaning that real stock returns are uncorrelated with expected and unexpected inflation (i.e. $\beta_1 = 0$, $\rho_{12} = 0$ and $\rho_{13} = 0$). According to the benchmark investor, real stock returns follow the iid process

$$r_{t+1} = \mu_1 + \varepsilon_{t+1}, \tag{5}$$

where (ε_t) is a series of independent normally distributed disturbances with mean zero and variance σ^2, uncorrelated with shocks in expected and unexpected inflation. An investor who believes that real stocks returns satisfy Equation (5) is not exposed to any inflation risk, thus forming a natural benchmark. Second, we consider the optimal asset allocation of a 'Fisherian' investor. This investor rules out a relation between real stock returns and expected inflation by setting $\beta_1 = 0$, but he does allow the innovations in stock returns to be correlated with shocks in expected and unexpected inflation ($\rho_{12} \neq 0, \rho_{13} \neq 0$). For this investor, stocks are a complete hedge against expected inflation, but not against unexpected inflation. Finally, we consider an 'agnostic' investor who bases his beliefs on the estimated VAR model of Equation (4). The difference in optimal stock allocations between the benchmark and the agnostic investor is interpreted as a measure for the exposure of stocks to inflation risk. This definition is motivated by the fact that the portfolio allocations of the two investors would be the same in a world without such risk.

2.2 Horizon effects

The inflation hedging literature has devoted considerable attention to horizon effects in the return-inflation relation; see e.g. Boudoukh and Richardson (1993), Solnik and Solnik (1997) and Schotman and Schweitzer (2000).

To gain insight into the risk-return trade-off in relation to the investment horizon, it is useful to derive the (conditional) mean and variance of the real stock returns in the agnostic investor's VAR model. Standard VAR calculations yield

$$\mathbb{E}_t[r_t(k)] = k\mu_1 + \beta_1[A(\beta_2, k)\pi_t^e + B(\beta_2, k)\mu_2], \tag{6}$$

$$\mathbb{V}\text{ar}_t[r_t(k)] = k\sigma_1^2 + 2B(\beta_2, k)\beta_1\sigma_{12} + C(\beta_2, k)\beta_1^2\sigma_2^2. \tag{7}$$

Here

$$A(\beta_2, k) = \sum_{i=1}^{k} \beta_2^{i-1}, \quad B(\beta_2, k) = \sum_{i=1}^{k}\sum_{j=1}^{i-1} \beta_2^{j-1}, \quad C(\beta_2, k) = \sum_{i=1}^{k}\Big(\sum_{j=1}^{i-1} \beta_2^{j-1}\Big)^2. \tag{8}$$

In Appendix A we show that, under certain conditions, the optimal share of wealth invested in the stock by a power utility investor is an increasing function of the expected real stock return and a decreasing function of the variance. The benchmark investor believes that real returns are iid. Equation (A.8) makes clear that there are no horizon effects for this investor, as both the mean and the variance grow linearly with the investment horizon.

Unless $\beta_1 = 0$, the real value of the agnostic investor's investment at $t + k$ will depend on the evolution of the inflation process, which induces horizon effects. From Equations (6) it becomes clear that the initial level of expected inflation affects the agnostic investor's optimal stock holdings. With $\beta_1 < 0$, high (low) initial expected inflation predicts low (high) expected returns and thus low (high) future stock returns, decreasing (increasing) the invested share in stocks. More formally, for the agnostic investor the conditional mean in Equation (6) is lower than $k\mathbb{E}_t[r_{t+1}] = k(\mu_1 + \beta_1\pi_t^e)$ for $\pi_t^e > \mathbb{E}[\pi_t] = \mu_2/(1 - \beta_2)$.

Predictability of stock returns from inflation rates may give rise to negative autocorrelation in stock returns. With $\beta_1\sigma_{12}$ sufficiently negative and $\sigma_{12} < 0$, a positive shock in stock returns generally coincides with a contemporaneous negative shock in expected inflation. Since $\beta_1 > 0$, the negative inflationary shock will decrease future stock returns. If the persistence in the inflation process is high, stocks will remain low until inflation rates have reached normal values again. Hence, the initial increase in stock returns is followed by a decrease, resulting in mean reversion. A similar mean reversion effect may occur for $\beta_1 < 0$ and $\sigma_{12} > 0$. In both cases the negative autocorrelation in stock returns causes the k-period return to be lower than $k\sigma_1^2$.

If the conditional variance of the k-period return is lower than $k\sigma_1^2$, the VAR investor considers stocks to be less risky in the long run. Consequently, he will allocate more wealth to stocks at longer investment horizons. The conditional variance in Equation (7) is lower than $k\mathbb{V}\text{ar}\,[r_{t+1}] = k\sigma_1^2$ for $\beta_1\sigma_{12}$ sufficiently negative. With $\beta_1\sigma_{12}$ not negative enough, the conditional variance may exceed $k\mu_1$. In this case the VAR investor considers stocks to be more risky in the long run, implying that he will allocate less wealth to stocks at longer investment horizons.

3 Bayesian approach

As emphasized in the introduction, accurate parameter estimates for the return equation of Model (4) are generally difficult to obtain. In particular, estimates for β_1 are usually characterized by large standard errors. This is due to the fact that, at short horizons, the time series properties of asset returns, which are highly volatile, differ considerably from those of the inflation process, which tends to be slowly moving and persistent. Nevertheless, it has been demonstrated in the asset allocation literature that a statistically significant relation in this type of regressions is not a necessary condition for economically significant results (see e.g. Kandel and Stambaugh, 1996). Barberis (2000) sketches three alternative ways to deal with return regressions that are characterized by low significance levels. The first option is to assume that insignificant coefficients are equal to zero. The second option is to ignore the considerable uncertainty in the estimated coefficients and to treat them as if they were exactly known. The third option is to account for parameter uncertainty. The latter option can be implemented by adopting a Bayesian approach for solving the investor's optimization problem as discussed in Section 2. Suppose that at time $t = T$, the investor estimates the parameters of Equation (4) using all available information about real returns and inflation. The estimated model parameters $\hat{\theta}$ and the information set $\mathcal{I}_T$ determine the conditional k-period return distribution with density function $p(r_T(k)|\mathcal{I}_T, \hat{\theta})$. For an investor who treats the estimated parameters as fixed, the optimization problem boils down to

$$E_T\left[u(W_{T+k})\right] = \max_{\lambda} \int u(W_{T+k})p(r_T(k)|\mathcal{I}_T, \widehat{\theta})dr_T(k). \tag{9}$$

Instead of using fixed parameter values, the Bayesian approach applies a posterior distribution $p(\theta|\mathcal{I}_T)$ to summarize the uncertainty about the parameters given the information set $\mathcal{I}_T$. This posterior distribution is used to weigh the conditional return distributions $p(r_T(k)|\mathcal{I}_T, \theta)$ in an objective function of the form

$$\max_{\lambda} \int \int u(W_{T+k}) p(r_T(k)|\mathcal{I}_T, \theta) p(\theta|\mathcal{I}_T) dr_T(k) d\theta. \tag{10}$$

This integral can be evaluated by means of simulation. A large number, say N, of end-of-period returns $r_T(k)$ is simulated by repeatedly drawing a set of model parameters from the posterior distribution, subsequently drawing the corresponding return value from the conditional distribution $p(r_T(k)|\mathcal{I}_T, \theta)$. The corresponding utility levels are then averaged over the N outcomes and the optimal value of λ is obtained using a numerical optimization routine. The details of this procedure are explained in the following two subsections.

3.1 Benchmark investor

The simple iid model of our benchmark investor involves parameter uncertainty about the mean and variance of stock returns; see Equation (5). Assuming normality of the errors $(\varepsilon_t)_t$, we can apply conventional methods (Zellner, 1971; Barberis, 2000) to obtain the posterior distribution $p(\mu, \sigma^2 \mid \mathcal{I}_T)$. Using an uninformative prior $p(\mu, \sigma^2) \propto 1/\sigma^2$, we first sample a value of σ^2 from an inverse gamma distribution with parameters $(T-1)/2$ and $(1/2)\sum_{t=1}^{T}(r_t - \overline{r})^2$. Subsequently, the corresponding value of μ can be sampled from a normal distribution with parameters $\overline{r}$ and σ^2/T. For each drawn pair of parameters (μ, σ^2), the corresponding conditional k-period real stock return, as seen from the perspective of the benchmark investor, can be sampled from a normal distribution with mean $k\mu$ and variance $k\sigma^2$. By repeating this procedure N times, we can approximate the posterior distribution $p(\mu, \sigma^2|\mathcal{I}_T)$. This yields a sample $r_T(k)^{(1)}, r_T(k)^{(2)} \ldots r_T(k)^{(N)}$ from the predictive distribution of k-period returns. The integral in Equation (10) is approximated by

$$\frac{1}{N}\sum_{i=1}^{N} u\big[W_{T+k}(r_T(k)^{(i)})\big]. \tag{11}$$

3.2 Agnostic investor

In contrast to the benchmark investor, the Fisherian and agnostic investors take into account the relation between real returns and inflation and estimate the VAR model of Equation (4). A recent overview of Bayesian estimation methods for such models is provided by Koop and Korobilis (2009). By defining $y_t = (r_{t+1}, \pi^e_{t+1}, \pi^u_{t+1})'$, $\theta = (\mu_1, \beta_1, \mu_2, \beta_2)'$ and

$$Z_t = \begin{pmatrix} 1 & \pi_t^e & 0 & 0 \\ 0 & 0 & 1 & \pi_t^e \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{12}$$

We can write the agnostic investor's VAR model as $\mathbf{y}_t = Z_t\theta + \varepsilon_t$, where the disturbance vector $\varepsilon_t = (\varepsilon_{1,t}, \varepsilon_{2,t}, \varepsilon_{3,t})$ is assumed to be independent multivariate normally distributed, with mean zero and covariance matrix Σ. Stacking the observations for all time periods, we write $\mathbf{y} = Z\theta + \varepsilon$. Here $\mathbf{y}$ and ε are $(3k \times 1)$ vectors and Z is a $(3k \times 4)$ matrix. With this setup we can use an uninformative independent normal-Wishart prior and conditional posterior distributions $p(\theta \mid \mathbf{y}, \Sigma^{-1}) \sim N(\overline{\theta}, \overline{V}_\theta)$ and $p(\Sigma^{-1} \mid \mathbf{y}, \theta) \sim W(\bar{S}^{-1}, \overline{v})$. Here

$$\overline{V}_\theta = \left(\sum_{t=1}^{T} Z_t'\Sigma^{-1}Z_t\right)^{-1}, \quad \overline{\theta} = \overline{V}_\theta \sum_{t=1}^{T} Z_t'\Sigma^{-1}\mathbf{y}_t; \tag{13}$$

with

$$\overline{v} = k, \quad \overline{S} = \sum_{t=1}^{T} (\mathbf{y}_t - Z_t\theta)(\mathbf{y}_t - Z_t\theta)'. \tag{14}$$

A Gibbs sampling algorithm is then used to draw sequentially from $p(\theta \mid \mathbf{y}, \Sigma^{-1})$ and $p(\Sigma^{-1} \mid \mathbf{y}, \theta)$. We then exploit the fact that given parameters (θ, Σ) the distribution of the k-period return is normal to obtain a sample from the predictive distribution $r_T^{(1)}, \ldots, r_T^{(N)}$, which is then used in Equation (11) to obtain expected utility for different stock allocations λ_t.

For the Fisherian investor, who believes that stocks are a hedge against expected but not against unexpected inflation, we adopt an approach similar to the one outlined here. We estimate the VAR model of Equation (4), while imposing the restriction $\beta_1 = 0$. We also impose this restrictions on the matrix Z_t in Equation (12).

4 Data

To obtain optimal portfolios using the methods discussed in Section 3, we need data on real stock returns, (proxies of) expected and unexpected inflation and a risk-free real rate. We focus on the United States and take the S&P 500 Total Return Index as our stock index. We use the data from the *Survey of Professional Forecasters* as a proxy for expected inflation and we also take realized inflation from this source.[1] Together, this yields a proxy for unexpected inflation. Furthermore, we use total inflation to convert nominal stock returns to real returns. For the risk-free real rate we use the real yield curve as provided by the U.S. Department of the Treasury, with maturities equal to five, seven and ten years.[2]

4.1 Sample period, data frequency and risk aversion

Although the economic literature has shown that it is reasonable to model inflation as a mean-reverting process, both the average level of inflation and the volatility of the inflation process differ considerably over subperiods. The differences between the Great Moderation, starting in the mid-1980's, and the previous inflationary period are particularly large (see Stock and Watson, 2005). To avoid any structural breaks, our sample period starts in the first quarter of 1985 and runs until the first quarter of 2010. In previous studies various data frequencies have been used. Fama and Schwert (1977) analyze monthly, quarterly and semi-annual data. Boudoukh and Richardson (1993) and Bekaert and Wang (2010) estimate long-term models using one-year to five-year (overlapping) stock returns. Given our relatively short sample period, we opt for quarterly (non-overlapping) data. The final input required for estimating our model is the coefficient of relative risk aversion ϕ. In a recent review article, Meyer and Meyer (2005) compare and synchronize the

[1] From 1991 onwards the Survey of Professional Forecasters contains real-time data about the realized inflation rate, which are subject to subsequent revisions. We do not make use of this additional data and simply use the 'final' value in our calculations. Our main reason for doing so is that the real-time data is not available for the 1985 – 1991 period, which is part of our sample period. The realized inflation rates available in the Survey of Professional Forecasters can be derived from the CPI series named 'USCONPRCE' taken from Thomson Reuters Datastream, corresponding to U.S. all urban seasonally adjusted CPI, provided by the Bureau of Labor Statistics.

[2] See http://www.ustreas.gov/offices/domestic-finance/debt-management/interest-rate/real_yield_historical.shtml.

empirical evidence obtained thus far in the literature. Based on studies by Friend and Blume (1975) and Blake (1996), they report risk aversion coefficients for wealth outcome variables between 2 and 5. In line with Guidolin and Timmermann (2007), we use a risk aversion coefficient $\phi = 5$ for our main analysis and provide additional results for a wide range of other values as a robustness check.

4.2 Expected inflation and stock index returns

In a recent study, Ang et al. (2007) show that surveys provide the best out-of-sample inflation forecasts. We use the one-quarter ahead inflation forecasts as available from the Survey of Professional Forecasters as a proxy for expected inflation. The deadline for forecast submission is typically in the second month of each quarter.[3] Forecasters are asked to predict the average quarter-to-quarter annualized inflation rate. To match the inflation forecasts with the appropriate stock returns, we observe that during our sample period the average quarterly inflation rate is highly correlated with the inflation rate obtained from dividing the mid-quarter CPI levels (the correlation during our sample period equals 0.98). Therefore, we associate to each quarterly inflation forecast the return on the stock index from the 15th of the second month of the quarter until the 15th of the second month in the next quarter.[4] As noticed by Ang et al. (2007), expectations of simple inflation rates differ from expectations of continuously compounded rates by Jensen's inequality term. Since this effect will generally be of little influence, we treat the forecasts of the simple rate inflation rate as forecasts of the continuously compounded inflation rates. Unexpected inflation rates are obtained by subtracting expected inflation from realized inflation rates.

4.3 Sample statistics

During our sample period the average quarterly real returns on the S&P 500 Total Return Index equals 1.7%, with standard deviation 7.8%; see Table 1. The inflation rate has a quarterly average value of 0.73%, with a standard error equal to 0.57%. Forecasted quarterly inflation, our proxy for expected inflation, equals on average 0.73% with standard error 0.24%. The difference between

[3] The deadline was generally the third week of the second month of the quarter during the period 1990 – 1998, the end of the second week in the period 1999 – 2004, and the middle or the start of the second week thereafter.

[4] Hence, the first quarter of the year starts on the 15th of February and runs until the 15th of May.

realized and forecasted inflation, our proxy for unexpected inflation, is on average 0.00% with standard deviation 0.48%. The risk-free real rate depends on the starting date and the maturity. Table 3 displays the real rate for various starting dates (the 15th of February of 2003 up to 2010) during our sample period and maturities of five, seven and ten years. For example, at the end date of our sample, the 15th of February 2010, the quarterly real yields equal 0.01%, 0.22% and 0.36% for maturities of five, seven and ten years, respectively.

Figure 1 displays quarterly stock index returns, together with expected and unexpected inflation rates over time. The different time series properties of stock returns and inflation rates (as mentioned in Section 1) become apparent immediately, as the return series is very volatile in comparison to the slowly moving processes of expected and unexpected inflation.

5 Empirical analysis

In this section we present the posterior distributions of the model parameters of the benchmark, agnostic and Fisherian investors. We discuss the estimation results and relate them to existing studies in the inflation hedging literature. Subsequently, we move on to the discussion of the optimal stock holdings of the three investors. The comparison of their optimal asset allocations allows us to quantify the impact of inflation risk on portfolio choice and to assess the relation between inflation risk and the investment horizon.

5.1 Setup

We draw samples of size $N = 1,000,000$ from the posterior parameter distributions corresponding to the models of Equations (4) and (5). We also estimate a restricted version of Equation (4), corresponding to a Fisherian investor who precludes a relation between real stock returns and expected inflation ($\beta_1 = 0$), but who does allow the innovations in stock returns to be correlated with shocks in expected and unexpected inflation ($\rho_{12} \neq 0, \rho_{13} \neq 0$). The means and standard deviations of the parameters of the three posterior distributions are displayed in Table 2.[5] To assess parameter stability and to illustrate the structural break in the data in the year 2009, we consider

[5] To save space, we do not display the full posterior distribution. It is available from the authors upon request.

different end dates for our sample period. The end dates run run from the first quarter of 2003 until the first quarter of 2010.

5.2 Posterior distributions

We first consider the estimation results for the benchmark investor in Table 2. We see that the expected quarterly real return is around 2% for the samples ending in 2003 – 2007, after which it decreases sharply. This decrease is accompanied by a considerable increase in real return volatility.

Turning to the other two investors, our main interest goes to the parameters β_1 (the Fisher coefficient), ρ_{12} and ρ_{13}. For all samples that end before 2009, the posterior means and standard errors of the model coefficient are fairly constant over time. In particular, the signs of β_1 (positive), ρ_{12} (negative), ρ_{13} (negative), and ρ_{23} (positive) are the same, regardless of the sample period. Hence, expected inflation positively affects expected stock returns ($\beta_1 > 0$), but return innovations are negatively correlated with unexpected inflation ($\rho_{13} < 0$). Furthermore, return innovations are negatively correlated with unexpected changes in expected inflation ($\rho_{12} < 0$) and changes in unexpected inflation are positively correlated with unexpected changes in expected inflation ($\rho_{23} > 0$). Interestingly, the relationship between stock returns and unexpected inflation changes considerably after 2008. With the end date of the sample period set to either 2009 or 2010, the signs of ρ_{12} and ρ_{13} turn out positive, but with relatively large standard errors. This change of signs is likely to be caused by the financial crisis. The stock market collapse was accompanied by a sudden, unexpected decrease in the inflation rate. The coefficient β_1 requires careful interpretation. The large standard deviations in Table 2 illustrate the magnitude of the parameter uncertainty problem. Regardless of the end date of the sample, the posterior mean of β_1 is positive, but the corresponding standard deviation is very large. In case of a classical VAR analysis, one should therefore seriously start questioning the usefulness of Equation (4). Also other parameters feature a lot of uncertainty. For example, with the end year set to 2003, the posterior mean of the intercept in Equation (5) equals 0.020, while its posterior standard deviation is more than three times as large. Similarly, the standard deviation of the μ_1-parameter is twice as large as its posterior mean in the VAR model of Equation (4).

As noticed in Section 3, the parameters β_1, ρ_{12} and ρ_{13} in the agnostic investor's model of

Equation (4) can be used to test the Fisher hypothesis. During all sample periods the standard deviation corresponding to the posterior distribution of β_1 is relatively large in comparison to the posterior mean. Hence, the approach of the benchmark and Fisher investors, who assume that $\beta_1 = 0$, does not seem unreasonable. Up to 2009, the posterior mean of ρ_{13} is negative (with a relatively small standard deviation), implying that stocks do not act as a complete hedge against unexpected inflation. As of 2009, the standard deviations corresponding to the posterior distributions of ρ_{12} and ρ_{13} are relatively large, and it is reasonable for the benchmark investor to assume that stocks are a complete hedge against both expected and unexpected inflation, in which case real stock returns are iid.[6]

Finally, we turn to a comparison of our results with existing studies. The estimated sign of β_1 differs widely across studies, see Schotman and Schweizer (2000) and Bekaert and Wang (2010). Despite the mixed evidence on the sign of β_1, all studies document a lot of uncertainty in the estimate of this coefficient. Moreover, there is some evidence that the sample period may affect the estimated sign of β_1. Several studies show that sustained periods of high inflation (and high expected inflation) can adversely affect real activity and lower stock returns; see e.g. Barnes et al. (1999). If a period of stagflation is included in the sample, it is likely that this effect will be captured. In a relatively stable inflationary environment, however, high expected inflation may reflect positive demand shocks, leading to higher company profits and stock returns. Notice that are sample period is restricted to the Great Moderation. The negative correlation between stock returns and unexpected inflation ($\rho_{13} < 0$) is in line with previous literature; see Schotman and Schweizer (2000).

5.3 Portfolio implications

We obtain optimal stock allocations for the period following the last day of our sample period, while assuming that stock returns and inflation rates continue to behave as during our sample period. As mentioned in Section 2.2, the initial level of expected inflation affects the optimal allocation to stocks for the agnostic investor via $\beta_1 > 0$. We initially abstract from this effect by setting

[6]An important issue to address is the possibility of a unit root in the autoregressive model for expected inflation. Fortunately, Sims et al. (1990) explain that unit roots are not a problem in a Bayesian setting.

the initial inflation rate equal to its long-term average value, so that it has no predictive power. We first consider the traditional approach of calculating optimal stock allocations and ignore parameter uncertainty. Table 4 displays the optimal stock holdings for our benchmark investor; see the column captioned 'benchmark (no PU)'. Portfolio weights are also provided for a Fisherian investor who believes that stocks are a complete hedge against expected inflation ($\beta_1 = 0$), but not against unexpected inflation ($\rho_{13} \neq 0$); see the column captioned 'Fisher (no PU)' in Table 4. Finally, Table 4 also reports the stock holdings for the investor who is agnostic about the relation between real returns and expected and unexpected inflation ('VAR (no PU)'). For all three investors we report the optimal stock allocations for different sample periods and investment horizons equal to five, seven and ten years.

5.3.1 Benchmark investor

For a benchmark investor who ignores parameter uncertainty the allocation to stocks decreases with the investment horizon. As shown by Barberis (2000), there are usually no horizon effects for such an investor, but this only holds if the real risk-free rate does not change with the maturity of the inflation-linked bond. As can be seen from Table 3, the risk-free rate increases with the maturity. Consequently, our inflation-linked bond becomes a more attractive investment opportunity in the long run, which is reflected in the decreasing share of stocks at longer investment horizons.

In Table 4 we observe certain differences in the benchmark investor's optimal stock allocations across sample periods. These differences are due to changes in the model parameters across sample periods, as well as to changes in the risk-free real rate over time. The changes in the model parameters result in changes in the mean and variance of stock index returns, affecting optimal portfolio holdings.

5.3.2 Agnostic investor

In addition to the term structure of real interest rates, another factor determines the existence and magnitude of horizon effects for the VAR investor. This factor is the predictability of stock returns on the basis of inflation. With $\beta_1 > 0$ and $\beta_1 \rho_{12}$ sufficiently negative, we would observe mean reversion in stock prices, making stocks less risky at longer investment horizons. With exception of the sample periods running until 2006 and 2007, the agnostic investor's stock allocations are *decreasing* in the investment horizon, despite the fact that $\beta_1\rho_{12} < 0$. Even with a constant term structure the stock allocations decrease with the investment horizon, which means that stock returns are not mean-reverting. As of 2009, the correlation between innovations in stock returns and shocks in expected inflation is positive ($\rho_{12} > 0$). In combination with $\beta_1 > 0$, this means that stock returns are mean averting, which makes them riskier at longer investment horizons.

The differences in the agnostic investor's optimal stock allocations across sample periods are due to changes in the risk-free real rate and the initial level of expected inflation. They are also caused by changes in the model parameters featuring Equation (4).

The differences in optimal portfolio holdings between the benchmark and the agnostic investor are substantial. Table 4 shows that the difference in stock allocations is particularly large for the samples ending in 2009 and 2010. For these samples the variance of stock returns is relatively high according to the agnostic investor's VAR model, due to the mean aversion in stock returns implied by $\rho_{12} > 0$ and $\beta_1 > 0$. This causes the agnostic investor's optimal stock allocation to be relatively low.

Since the risk-free real rate may depend on the level of expected inflation, it is more realistic to set the initial value of expected inflation equal to its value at the end date of our sample period (which coincides with the start date of our simulations). This procedure ensures that we take a realistic combination of the risk-free real rate and the level of expected inflation. We notice that the initial level of expected inflation only matters for the agnostic investor; in the VAR model of Equation (4) the initial value of expected inflation affects the eventual portfolio holdings. For all sample periods the initial level of expected inflation is below its long-term average value. Hence, the agnostic investor starts with a relatively low level of expected inflation, yielding relatively low expected real returns due to $\beta_1 > 0$. The agnostic investor faces an additional horizon effect in addition to the impact of the term structure of real interest rates and potential predictability effects. If the highly persistent process of expected inflation is below its long-term average value, it will slowly increase over time. Due to $\beta_1 > 0$, expected stock returns will also increase over time. This makes stocks a more attractive investment in the long run. Clearly, such a horizon effect is not present in the benchmark investor's stock holdings. Table 5 displays the allocation to stocks for the agnostic investor based on the more realistic initial levels of expected inflation (see the columns captioned 'no PU'). The differences in stock allocation between the benchmark and the agnostic investor are even larger than before and amount to as much as 40 percentage points for the sample ending in 2009 and a five-year investment horizon.

5.3.3 Fisherian investor

We now turn to our Fisherian investor, who believes that stocks are a complete hedge against expected, but not against unexpected inflation. Table 4 makes clear that his stock allocations do not differ much from those of the benchmark investor. This is due to the fact that the impact of unexpected inflation on real returns only occurs via the error term in the restricted version of the model in Equation (4), and not via the Fisher coefficient β_1. It is exactly the Fisher coefficient that plays an crucial role in the existence of predictability and mean reversion effects and the determination of the conditional mean and variance of real returns (see Equations (6) and (7)).

5.3.4 Parameter uncertainty

Until sofar we discussed optimal portfolio weights that did not yet account for parameter uncertainty. Table 4 also reports the optimal stock holdings based on the Bayesian approach. With parameter uncertainty, all three investor face an additional horizon effect. From the perspective of the benchmark investor who incorporates parameter uncertainty, returns are no longer independent but positively correlated. If returns are high in a given period, it is likely that the state of the world is one with a high realization of the parameter μ_1 in Equation (5), which implies that the returns will be high in subsequent periods as well. This positive correlation makes the variance of the multi-period real returns grow faster than linearly over time. Since parameter uncertainty increases with the investment horizon, it induces considerable horizon effects. As can be seen from Table 4, based on the sample ending in 2004 the benchmark investor allocates 17 percentage points less to stocks with a ten-year investment horizon, relative to a five-year investment horizon. For the agnostic investor the impact of parameter uncertainty is even larger. The agnostic investor's VAR model contains more parameters than the simple model of the benchmark investor and is therefore subject to higher degree of parameter uncertainty. Also the portfolio weights of the benchmark and the agnostic investor differ more substantially if parameter uncertainty is accounted for, according to Table 4. With a ten-year investment horizon, the difference in stock allocations amounts to as much as almost 20 percentage points for the sample ending in 2008. The agnostic investor allocates much less of his wealth to stocks than the benchmark investor. Table 5 provides optimal stock allocations for the agnostic investor based on an initial level of expected inflation equal to its value at the end date of our sample period, taking into account parameter uncertainty (see the columns captioned 'with PU'). For several sample periods the influence of the initial level of expected inflation on portfolio allocations is dominated by the impact of parameter uncertainty.

As explained before, the difference in optimal stock allocations between the benchmark and the agnostic investor provides an alternative measure for the exposure of stock returns to inflation risk that explicitly accounts for parameter uncertainty. This definition is motivated by the fact that the portfolio allocations of the two investors would be the same in a world free of inflation risk. The difference in portfolio weights between the benchmark and agnostic investors reveals a substantial exposure of stock returns to inflation risk.

6 Robustness analysis

To ensure that the empirical results found in Section 5 are robust against changes in the model specification, we run several robustness checks.

6.1 Level of risk aversion

Our results depend on the investor's coefficient of relative risk aversion. As discussed before, empirical evidence suggests that this coefficient should be between 2 and 5. Barberis (2000) obtains results for ϕ as high as 20. As an illustration, Table 6 reports the Bayesian stock allocations for the benchmark and the agnostic investor for $\phi = 2, 3, 4$ and $\phi = 10$. The main conclusion is that the impact of inflation risk can be very substantial for investors who are relatively little risk averse. This effect is caused primarily by the fact that such investors allocate a relatively large proportion of their wealth to stocks. The inflation exposure of stocks, as measured by the difference in optimal stock allocation between the benchmark investor and the agnostic investor, is of a significant magnitude for a wide range of plausible risk aversion coefficients.

6.2 Dividend yield as an additional predictive factor

As noted by Ang and Bekaert (2007), the 'conventional wisdom' in the financial literature is that dividend yields strongly predict excess returns, with stronger predictability at longer investment horizons. Ang and Bekaert (2007) critically re-examine this dogma using long data sets for four countries (United States, France, Germany, and United Kingdom). They find that dividend yields predict excess returns only at short horizons. In another critical study, Boudoukh et al. (2008) show that the use of overlapping returns in combination with highly persistent predictive variables (such as the dividend yield) results in estimated coefficients for the predictive variables that are almost perfectly correlated across horizons under the null hypothesis of no predictability.

Thanks to our Bayesian approach the controversy in the literature about the predictive power of the dividend yield does not have to refrain us from considering a fourth investor who includes dividend yields in his VAR model, in addition to expected and unexpected inflation. By analyzing the asset allocations of this additional investor, we can explore the role of expected and unexpected inflation in the situation that dividend yields are already part of the agnostic investor's VAR model. Even if dividend yields do not significantly affect stock returns, our Bayesian approach ensures that we take into account all information that is contained in the relation between stock returns and dividend yields.

6.2.1 Dividend yields and investor beliefs

Similar to Ang and Bekaert (2007), we focus on the one-year rolling window dividend to circumvent seasonality. That is, we aggregate the dividends paid in the four quarters prior to time t and divide this by the value of the stock index at time t, resulting in $D_t^4 = [D_t + D_{t-1} + D_{t-3} + D_{t-4}]/P_t$. We download the rolling-window dividend yields corresponding to the S&P 500 Total Return Index from Thomson Reuters Datastream. Its sample mean equals 2.44% during the 1985 – 2010 period, with a standard deviation of 0.90%; see the last column in Table 1. To obtain the asset allocations of our fourth investor, we proceed in a similar way as before. We specify a four-dimensional restricted VAR model in which the log dividend yield $d_t = \log(D_t^4)$ affects stock returns over the period from time t until $t+1$. The new VAR model is given by

$$\begin{pmatrix} r_{t+1} \\ \pi^e_{t+1} \\ \pi^u_{t+1} \\ d_{t+1} \end{pmatrix} = \begin{pmatrix} \mu_1 \\ \mu_2 \\ 0 \\ \mu_3 \end{pmatrix} + \begin{pmatrix} 0 & \beta_1 & 0 & \gamma \\ 0 & \beta_2 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \beta_3 \end{pmatrix} \begin{pmatrix} r_t \\ \pi^e_t \\ \pi^u_t \\ d_t \end{pmatrix} + \begin{pmatrix} \varepsilon_{1,t+1} \\ \varepsilon_{2,t+1} \\ \varepsilon_{3,t+1} \\ \varepsilon_{4,t+1} \end{pmatrix}, \tag{15}$$

where $(\varepsilon_{1,t}, \varepsilon_{2,t}, \varepsilon_{3,t}, \varepsilon_{4,t})$ is a series of independent multivariate normally distributed disturbances, with mean zero and covariance matrix $\mathbb{C}\text{ov}\,(\varepsilon_{1,t}, \varepsilon_{2,t}, \varepsilon_{3,t}, \varepsilon_{4,t}) = \Sigma$. The VAR model of Equation (15) is an extension of Equation (4). The log dividend yield calculated over the four quarters prior to time t affects stock returns over the period from t until $t+1$ in the first equation of the extended VAR model. The fourth equation of the new VAR model specifies the log dividend yield as a simple autoregressive process. We apply the Bayesian methods explained in Section 3 to obtain optimal stock allocations, taking into account parameter uncertainty.

For the purpose of comparison, we also estimate a simple two-dimensional VAR model for returns and dividend yields (similar to Barberis, 2000). This VAR model is specified as

$$\begin{pmatrix} r_{t+1} \\ d_{t+1} \end{pmatrix} = \begin{pmatrix} \eta_1 \\ \eta_2 \end{pmatrix} + \begin{pmatrix} 0 & \gamma_1 \\ 0 & \gamma_2 \end{pmatrix} \begin{pmatrix} r_t \\ d_t \end{pmatrix} + \begin{pmatrix} \nu_{1,t+1} \\ \nu_{2,t+1} \end{pmatrix}, \tag{16}$$

where $(\nu_{1,t}, \nu_{2,t})$ is a series of independent multivariate normally distributed disturbances, with mean zero and covariance matrix $\mathbb{C}\text{ov}\,(\nu_{1,t}, \nu_{2,t}) = \Omega$. The simple VAR model corresponds to an agnostic investor who ignores the role of expected and unexpected inflation.

6.2.2 Estimation results

The estimation results for the simple VAR model are given in the upper part of Table 7. Dividend yields positively affect expected stock returns ($\gamma_1 > 0$). The positive coefficient of the dividend yield in the return equation results in mean reversion in stock returns, since the correlation between return innovations and shocks in the dividend yield is highly negative ($\omega_{12} < 0$). Furthermore, from the last equation of the VAR model we observe that the log dividend yield is a highly persistent process ($\gamma_2 > 0.95$). The estimation results in Table 7 make clear that the parameters of the simple VAR model are relatively stable over time.

Also in the extended VAR model dividend yields positively affect expected stock returns ($\beta_3 > 0$). Moreover, expected inflation negatively affects expected stock returns for the samples ending before or in 2008 ($\beta_1 < 0$). Moreover, the correlation between innovations in stock returns and shocks in expected inflation is negative in these cases ($\rho_{12} < 0$). Also the correlation between innovations in stock returns and unexpected inflation is negative ($\rho_{13} < 0$). For the samples ending in 2009 or 2010, the coefficient of expected inflation (β_1) in the return equation is positive and the aforementioned correlations (ρ_{12} and ρ_{13}) are positive. All subsamples feature $\beta_1 \rho_{12} > 0$, which means that the inflation rate partly offsets the mean reversion effects induced by the dividend yield.

6.2.3 Portfolio implications

The optimal stock allocations for both the simple and the extended VAR model are reported in Table 8. The initial level of expected inflation and the dividend yield are set to their values at the end date of our sample period, which are all below average. As in Section 5, we face horizon effects due to the (1) term structure of real yields, (2) the initial level of the predictor variables (expected inflation and/or dividend yields) and (3) parameter uncertainty.

We start with the simple VAR model of Equation (16). With a flat term structure of real yields, no parameter uncertainty and the initial dividend yield equal to the sample average, the optimal allocation to stocks would increase with the investment horizon because of the mean reversion in stock returns (caused by the predictability of stock returns from the dividend yield). With the aforementioned factors causing horizon effects, we only observe increasing stock holdings for the samples ending in 2006 and 2007. Notice that in these years the term structure of real yields was relatively flat. For the other sample periods, the influence of the term structure of real yields and parameter uncertainty results in optimal weights that decrease with the investment horizon. Although the parameters featuring the simple VAR model are relatively stable over time, we observe substantial differences in its optimal stock holdings across sample periods. These differences are due to changes in the term structure of real yields over time and the use of different initial levels of the dividend yield.

The optimal stock holdings based on the simple VAR model, which ignores the role of expected and unexpected inflation, are substantially higher than the weights obtained from the extended VAR model of Equation (15). As explained in Section 6.2.2, the inflation rate partly offsets the mean reversion effects caused by the dividend yield. Consequently, the simple VAR model of Equation (16) understates real return volatility in comparison to the extended VAR model. Moreover, with $\beta_1 < 0$ and the initial level of expected inflation below average, stocks become less attractive at longer investment horizons in the extended VAR model. The simple VAR model ignores this horizon effect. The overall effect is that the latter model allocates too much wealth to stocks.

6.3 Unrestricted VAR models

The VAR models of Section 2 and 6.2.1 impose several a priori restrictions on the model coefficients. For example, the VAR models do not feature any autoregressive effects in the return equation. These restrictions facilitate the interpretation of the effect that shocks in (un)expected inflation have on stock returns. It is not difficult to adjust our approach to fully unrestricted VAR models. Since this leads to very similar empirical results as in Section 5, we do not report any estimation results here.[7]

7 Conclusions

A widely adopted view in the economic literature is that an asset is a good hedge against inflation if the Fisher hypothesis holds true; i.e. if nominal asset returns move one-to-one with expected and/or unexpected inflation. In a regression of volatile stock returns on slowly moving inflation rates the estimated Fisher coefficient is often characterized by large standard errors, due to which the Fisher hypothesis cannot be rejected.

It has been demonstrated in the asset allocation literature that a statistically significant relation in regressions of stock returns on predictor variables is not a necessary condition for economically significant results (see e.g. Kandel and Stambaugh, 1996). But investors who base their optimal

[7] These results are available from the authors upon request.

investment portfolios on the assumption that stocks are a good hedge against inflation ignore the information contained in the Fisher coefficient. This raises the question whether stocks are still viewed as a good hedge against inflation if we take into account the parameter uncertainty involved with the Fisher coefficient.

In our empirical study we indeed find little traditional statistical evidence against the Fisher hypothesis, suggesting that stocks are a good hedge against both expected and unexpected inflation. Very different results emerge if we adopt a Bayesian approach. We analyze the difference in stock allocations between an agnostic investor and a benchmark investor, where the latter assumes that the Fisher hypothesis holds true and the former does not. In a 'Fisherian' world free of inflation risk the benchmark and agnostic investors would have the same optimal portfolio weights. Hence, the difference in portfolio weights between the two investors provides a measure of the inflation risk exposure of stocks. By means of a Bayesian approach to portfolio optimization we explicitly take into account the parameter uncertainty involved with the Fisher coefficient and the other model coefficients. Our main finding is that the stock allocations of the benchmark and agnostic investors differ substantially, revealing a substantial exposure of stock returns to inflation risk.

Our results have important implications for short-term and long-term investors. Accurate modeling of the relation between stock returns and inflation is crucial to make optimal portfolio choices. Furthermore, instead of simply ignoring parameter uncertainty, this uncertainty can be used as an additional source of information, which strongly affects optimal asset allocations. Possible extensions of our analysis include a comparison across several assets, countries, and sample periods. We leave this as a topic for future research.

References

Ang, A. and Bekaert, G. (2002). International Asset Allocation With Regime Shifts. *Review of Financial Studies* 15, 1137-1187.

Ang, A. and Bekaert, G. (2007). Stock Return Predictability: Is it There? *Review of Financial Studies* 20, 651-707.

Ang, A. Bekaert, G., and Wei, M. (2007). Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better? *Journal of Monetary Economics* 54, 1163-1212.

Barberis, N. (2000). Investing for the Long Run When Returns are Predictable. *Journal of Finance* 55, 225-264.

Barnes, M., Boyd, J.H, and Smith B.D. (1999). Inflation and Asset Returns. *European Economic Review* 43, 37-754.

Bekaert, G. and Wang, X.S. (2010). Inflation Risk and the Inflation Risk Premium. *Economic Policy* 25, 755-806.

Blake, D. (1996). Efficiency, Risk Aversion and Portfolio Insurance: An Analysis of Financial Asset Portfolios Held by Investors in the United Kingdom. *Economic Journal* 106, 1175-1192.

Bodie, Z. (1976). Common Stocks as a Hedge against Inflation. *Journal of Finance* 31, 459-470.

Boudoukh, J. and Richardson, M. (1993). Stock Returns and Inflation: A Long-Horizon Perspective. *American Economic Review* 83, 1346-1355.

Boudoukh, J., Richardson, M., and Whitelaw, R.F (2008). The Myth of Long-Horizon Predictability. *Review of Financial Studies* 21, 1577-1605.

Campbell, J.Y., Chan, Y.L., and Viceira, L.M. (2003). A Multivariate Model of Strategic Asset Allocation. *Journal of Financial Economics*.

Fama, E.F. and MacBeth, J.D. (1974). Tests of the Multiperiod Two-Parameter Model. *Journal of Financial Economics* 1, 43-66.

Fama, E.F. and Schwert, G.W. (1977). Asset Returns and Inflation. *Journal of Financial Economics* 5, 115-146.

Fisher, I. (1930). *The Theory of Interest*. New York: MacMillan.

Friend, I. and Blume, M.E. (1975). The Demand for Risky Assets. *American Economic Review*

65, 900-922.

Guidolin, M. and Timmermann, A. (2007). Asset Allocation under Multivariate Regime Switching. *Journal of Economic Dynamics and Control* 31, 3503-3544.

Gültekin, N.B. (1983). Stock Market Returns and Inflation: Evidence from Other Countries. *Journal of Finance* 38, 49-65.

Jaffe, J.F. and Mandelker, G. (1976). The 'Fisher Effect' for Risky Assets: An Empirical Investigation. *Journal of Finance* 31, 447-458.

Kandel, S. and Stambaugh, R.F. (1996). On the Predictability of Stock Returns: An Asset-Allocation Perspective. *Journal of Finance* 51, 385-424.

Koop, G. and Korobilis, D. (2009). Bayesian Multivariate Time Series Methods for Empirical Macroeconomics. MPRA Paper 20125.

Lintner, J. (1973). *Inflation and Common Stock Prices in a Cyclical Context*. National Bureau of Economic Research 53rd Annual Report, New York.

Meyer, D.J. and Meyer, J. (2005). Relative Risk Aversion: What Do We Know? *Journal of Risk and Uncertainty* 31, 243-262.

Nelson, C.R. (1976). Inflation and Rates of Return on Common Stock. *Journal of Finance* 31, 471-483.

Roll, R. (1972). Interest Rates on Monetary Assets and Commodity Price Index Changes. *Journal of Finance* 27, 251-277.

Schotman, P.C. and Schweitzer, M. (2000). Horizon Sensitivity of the Inflation Hedge of Stocks. *Journal of Empirical Finance* 7, 301-315.

Sims, C.A., Stock, J.H., Watson, M.W. (1990). Inference in Linear Time Series Models with Some Unit Roots. *Econometrica* 58, 16182.

Solnik, B. (1983). The Relation between Stock Prices and Inflationary Expectations: The International Evidence. *Journal of Finance* 38, 35-48.

Solnik, B. and Solnik, V. (1997). A Multi-Country Test of the Fisher Model for Stock Returns. *Journal of International Financial Markets, Institutions and Money* 7, 289-301.

Stock, J.H. and Watson, M.W. (2006). Why has U.S. Inflation Become Harder to Forecast? NBER Working Paper 12324.

Zellner, A. (1971). *An Introduction to Bayesian Inference in Econometrics*, New York.

Wakker, P.P. (2008). Explaining the Characteristics of the Power (CRRA) Utility Family. *Health Economics* 17, 1329-1344.

Figure 1: Quarterly real returns and expected and unexpected inflation

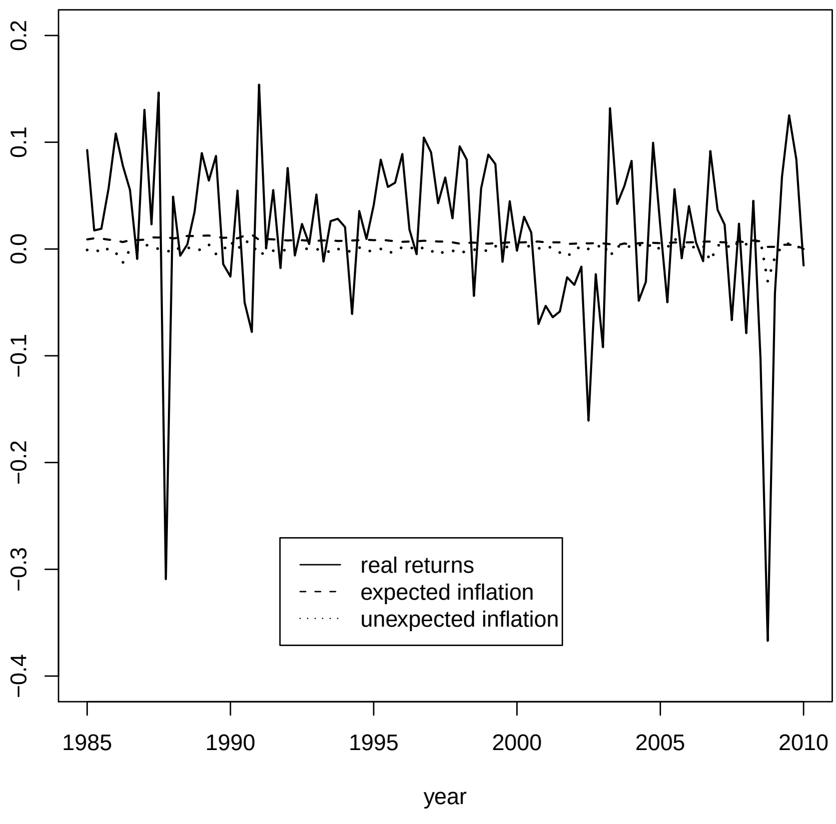

Table 1: Sample statistics for stock returns, inflation and dividend yields

	returns	**exp. infl.**	**unexp. infl.**	**total infl.**	**div. yield**
mean	1.72	0.73	0.00	0.73	2.44
median	2.38	0.70	0.00	0.76	2.20
std. dev.	7.79	0.24	0.48	0.57	0.90
skewness	-1.94	0.19	-2.58	-2.40	0.40
excess kurtosis	7.48	0.52	14.79	14.16	-0.93
5% quantile	-7.88	0.38	-0.68	-0.22	1.22
10% quantile	-6.11	0.49	-0.37	0.29	1.34
90% quantile	9.06	1.08	0.50	1.28	3.60
95% quantile	10.90	1.21	0.71	1.52	3.75
99.5% quantile	15.03	1.30	0.89	1.93	3.96

This table displays sample statistics for quarterly stock returns, expected inflation, unexpected inflation and total inflation (all measured in %), as well as one-year rolling-window dividend yields in % (see Section 6) during the period 1985 – 2010.

Table 2: Means and standard deviations of the posterior parameter distributions

	benchmark		**Fisher**		**VAR**		**benchmark**		**Fisher**		**VAR**	
	mean	**std.dev.**	**mean**	**std.dev.**	**mean**	**std.dev.**	**mean**	**std.dev.**	**mean**	**std.dev.**	**mean**	**std.dev.**
2003							**2007**					
μ_1	0.0199	0.0084	0.0181	0.0082	-0.0162	0.0322	0.0222	0.0073	0.0220	0.0071	0.0084	0.0258
β_1					4.2262	3.8395					1.7645	3.2368
μ_2			0.0005	0.0004	0.0006	0.0004			0.0004	0.0003	0.0005	0.0003
β_2			0.9334	0.0423	0.9250	0.0432			0.9409	0.0388	0.9361	0.0400
ρ_{12}			-0.2701	0.1104	-0.2731	0.1102			-0.2925	0.0982	-0.2934	0.0983
ρ_{13}			-0.2277	0.1130	-0.2242	0.1131			-0.2277	0.1019	-0.2170	0.1026
ρ_{23}			0.5618	0.0817	0.5618	0.0820			0.4338	0.0871	0.4335	0.0872
σ_1^2	0.0051	0.0009	0.0052	0.0009	0.0051	0.0009	0.0047	0.0007	0.0047	0.0007	0.0047	0.0007
σ_2^2			0.0009	0.0002	0.0009	0.0002			0.0008	0.0001	0.0008	0.0001
σ_3^2			0.0106	0.0019	0.0106	0.0019			0.0131	0.0020	0.0131	0.0020
2004							**2008**					
μ_1	0.0230	0.0082	0.0212	0.0079	0.0085	0.0297	0.0201	0.0071	0.0208	0.0069	0.0066	0.0254
β_1					1.6204	3.6023					1.8449	3.2022
μ_2			0.0004	0.0004	0.0005	0.0004			0.0004	0.0003	0.0004	0.0003
β_2			0.9433	0.0425	0.9394	0.0436			0.9381	0.0406	0.9350	0.0413
ρ_{12}			-0.2858	0.1064	-0.2860	0.1065			-0.2251	0.0995	-0.2263	0.0993
ρ_{13}			-0.2353	0.1089	-0.2325	0.1093			-0.2315	0.0997	-0.2173	0.1001
ρ_{23}			0.4931	0.0875	0.4926	0.0877			0.4178	0.0864	0.4177	0.0870
σ_1^2	0.0051	0.0008	0.0051	0.0009	0.0051	0.0009	0.0047	0.0007	0.0047	0.0007	0.0047	0.0007
σ_2^2			0.0009	0.0002	0.0009	0.0002			0.0009	0.0001	0.0009	0.0001
σ_3^2			0.0110	0.0019	0.0110	0.0019			0.0137	0.0021	0.0137	0.0021
2005							**2009**					
μ_1	0.0224	0.0079	0.0214	0.0077	0.0073	0.0279	0.0145	0.0081	0.0147	0.0079	-0.0149	0.0281
β_1					1.8116	3.4386					3.9422	3.5832
μ_2			0.0004	0.0003	0.0005	0.0003			0.0005	0.0003	0.0005	0.0003
β_2			0.9410	0.0401	0.9363	0.0414			0.9279	0.0404	0.9283	0.0406
ρ_{12}			-0.2967	0.1030	-0.2986	0.1027			0.1070	0.1013	0.1074	0.1015
ρ_{13}			-0.2226	0.1074	-0.2151	0.1075			0.1663	0.1000	0.1741	0.0996
ρ_{23}			0.4947	0.0854	0.4947	0.0853			0.5916	0.0670	0.5917	0.0672
σ_1^2	0.0050	0.0008	0.0050	0.0008	0.0050	0.0008	0.0062	0.0009	0.0063	0.0009	0.0063	0.0009
σ_2^2			0.0009	0.0002	0.0009	0.0002			0.0012	0.0002	0.0012	0.0002
σ_3^2			0.0110	0.0018	0.0110	0.0018			0.0248	0.0037	0.0248	0.0037
2006							**2010**					
μ_1	0.0218	0.0076	0.0217	0.0074	0.0071	0.0269	0.0165	0.0079	0.0165	0.0077	-0.0003	0.0259
β_1					1.8806	3.3405					2.2814	3.3453
μ_2			0.0004	0.0003	0.0004	0.0003			0.0003	0.0003	0.0003	0.0003
β_2			0.9456	0.0391	0.9408	0.0404			0.9523	0.0416	0.9524	0.0418
ρ_{12}			-0.2920	0.1006	-0.2937	0.1005			0.1122	0.0992	0.1120	0.0994
ρ_{13}			-0.2017	0.1054	-0.1904	0.1058			0.1861	0.0971	0.1912	0.0969
ρ_{23}			0.4716	0.0858	0.4714	0.0857			0.5475	0.0707	0.5475	0.0707
σ_1^2	0.0048	0.0008	0.0049	0.0008	0.0048	0.0008	0.0062	0.0009	0.0063	0.0009	0.0063	0.0009
σ_2^2			0.0009	0.0001	0.0009	0.0001			0.0013	0.0002	0.0013	0.0002
σ_3^2			0.0119	0.0019	0.0118	0.0019			0.0242	0.0035	0.0242	0.0035

This table displays the means and standard deviations of the posterior distributions for the parameters of the models in Equations (5) and (4). The parameters σ_2^2 and σ_3^2 have been multiplied by a factor 1,000. The models in this table correspond to three investors: (1) a benchmark investor who believes that stocks are a complete hedge against expected and unexpected inflation (see the column captioned 'benchmark'), (2) a Fisherian investor who believes that stocks are only a complete hedge against expected inflation ('Fisher'), and (3) an agnostic investor who allows real stocks returns to depend on both expected and unexpected inflation ('VAR'). Estimation results are provided for quarterly samples starting in 1985 and ending in the first quarter of the years 2003 up to 2010, as indicated in the first column.

Table 3: Term structure of real interest rates (in % per quarter)

	maturity (years)		
start	**5**	**7**	**10**
2003	0.32	0.44	0.50
2004	0.21	0.32	0.43
2005	0.26	0.32	0.40
2006	0.52	0.52	0.53
2007	0.59	0.60	0.59
2008	0.17	0.29	0.37
2009	0.30	0.35	0.42
2010	0.10	0.22	0.36

This table displays the real yield as provided by the U.S. Department of the Treasury. The starting date of the real yield is the 15th of February of the year given in the first column.

Table 4: Optimal allocation to stocks (in %) for different investors and various investment horizons

	benchmark (no PU)	**benchmark (with PU)**	**Fisher (no PU)**	**Fisher (with PU)**	**VAR (no PU)**	**VAR (with PU)**
2003						
5	75.5	61.1	68.0	55.3	70.5	51.8
7	71.0	53.7	63.0	48.8	63.0	42.5
10	68.8	47.2	60.5	42.8	59.0	33.9
2004						
5	92.4	75.1	85.5	70.0	94.0	63.8
7	88.4	67.2	81.0	62.3	90.0	52.5
10	84.3	58.3	77.5	54.3	87.0	40.4
2005						
5	89.2	73.2	85.5	71.2	92.5	65.4
7	87.0	66.8	83.0	64.4	90.5	56.0
10	84.3	59.1	81.0	56.7	89.5	44.6
2006						
5	78.9	65.3	79.0	65.1	76.0	56.5
7	78.8	61.2	78.0	61.6	75.5	49.7
10	78.8	56.1	79.0	56.5	76.5	41.1
2007						
5	79.7	66.4	79.0	65.8	85.5	62.6
7	79.6	62.3	78.5	62.4	85.5	55.9
10	80.1	57.5	80.0	57.9	88.0	46.7
2008						
5	89.0	74.5	92.5	77.2	84.0	62.4
7	84.2	66.6	87.0	68.9	78.0	50.6
10	80.9	59.0	85.0	61.7	75.0	40.8
2009						
5	46.2	39.8	47.0	40.7	35.0	26.5
7	44.6	36.5	45.0	37.4	31.0	22.4
10	42.4	32.2	42.5	33.0	28.0	18.2
2010						
5	59.9	51.4	60.0	51.2	47.5	32.7
7	56.0	45.6	55.5	45.8	41.5	25.7
10	51.3	38.9	50.5	39.1	35.5	19.0

This table displays the optimal stock allocations (in % of initial real-term wealth) for (1) a benchmark investor who believes that stocks are a complete hedge against expected and unexpected inflation (see the column captioned ‘benchmark’), (2) a Fisherian investor who believes that stocks are only a complete hedge against expected inflation (‘Fisher’), and (3) an agnostic investor who allows real stocks returns to depend on both expected and unexpected inflation (‘VAR’). We consider optimal stock allocations that account for parameter uncertainty (‘with PU’) and allocations that do not (‘no PU’). The investment horizons are five, seven and ten years, as indicated in the first column of the table. The allocations correspond are based on quarterly samples starting in 1985 and ending in the first quarter of the years 2003 up to 2010. The initial level of expected inflation is set to its long-term average value.

Table 5: The agnostic investor's optimal allocation to stocks (in %) for different investment horizons

	no PU	with PU		no PU	with PU
2003			**2007**		
5	46.5	32.9		81.5	59.8
7	44.5	29.0		82.0	53.7
10	45.5	25.4		85.0	45.3
2004			**2008**		
5	81.5	52.1		89.5	68.3
7	79.0	43.4		82.5	55.4
10	78.0	34.4		79.0	44.1
2005			**2009**		
5	83.0	57.3		7.0	6.1
7	82.0	50.0		10.0	7.9
10	82.0	40.6		12.5	8.6
2006			**2010**		
5	72.5	53.1		36.5	22.4
7	72.5	47.0		32.5	18.3
10	74.0	39.2		28.5	14.3

This table displays the optimal stock allocations (in % of initial real-term wealth) for an agnostic investor who assumes that real stocks returns depend on both expected and unexpected inflation. We consider allocations that account for parameter uncertainty ('with PU') and allocations that do not ('no PU'). The investment horizons are five, seven and ten years, as indicated in the first column. The allocations are based on quarterly samples starting in 1985 and ending in the first quarter of the years 2003 up to 2010. The initial level of expected inflation is set to its value at the end of the sample period.

Table 6: Optimal stock allocations (in %) for different values of the risk aversion parameter

	benchmark $\phi = 2$	VAR $\phi = 2$	benchmark $\phi = 3$	VAR $\phi = 3$	benchmark $\phi = 4$	VAR $\phi = 4$	benchmark $\phi = 10$	VAR $\phi = 10$
2003								
5	100.0	82.6	100.0	64.5	76.5	41.4	29.9	16.2
7	100.0	73.6	89.3	57.4	67.5	36.5	26.0	14.2
10	100.0	65.9	79.7	49.3	59.7	32.1	22.7	12.3
2004								
5	100.0	100.0	100.0	85.8	93.2	65.2	36.9	25.6
7	100.0	97.1	100.0	72.0	83.7	54.4	32.7	21.2
10	100.0	83.4	96.2	58.1	73.2	43.4	28.1	16.7
2005								
5	100.0	100.0	100.0	94.6	91.0	71.7	36.1	28.2
7	100.0	100.0	100.0	82.5	83.4	62.7	32.7	24.5
10	100.0	93.5	97.2	67.9	74.3	51.1	28.6	19.7
2006								
5	100.0	100.0	100.0	88.2	81.9	66.6	32.2	26.1
7	100.0	100.0	100.0	78.2	76.9	59.0	29.9	23.0
10	100.0	91.4	93.5	65.7	70.7	49.4	27.1	19.1
2007								
5	100.0	100.0	100.0	98.7	83.1	74.9	32.8	29.5
7	100.0	100.0	100.0	88.6	78.2	67.2	30.6	26.3
10	100.0	98.0	95.6	75.0	72.5	56.9	28.0	22.1
2008								
5	100.0	100.0	100.0	100.0	92.8	85.1	36.8	33.7
7	100.0	100.0	100.0	89.0	83.3	69.0	32.6	27.4
10	100.0	94.8	97.2	72.4	74.2	55.2	28.6	21.6
2009								
5	99.9	15.5	67.2	10.2	50.1	7.6	19.5	3.0
7	92.8	20.5	62.1	13.4	46.0	9.9	17.7	3.9
10	84.1	23.3	55.5	14.9	40.9	11.0	15.5	4.2
2010								
5	100.0	57.4	85.8	38.0	64.5	28.2	25.1	11.0
7	100.0	48.0	76.9	31.3	57.4	23.1	22.1	8.9
10	98.2	38.9	66.6	24.9	49.3	18.2	18.7	6.9

This table displays the optimal stock allocations (in % of initial real-term wealth) for the benchmark and the agnostic investor, for different levels of risk aversion ϕ. The optimal stock allocations account for parameter uncertainty. The investment horizons are five, seven and ten years, as indicated in the first column. The allocations are based on quarterly samples starting in 1985 and ending in the first quarter of the years 2003 up to 2010, as also indicated in the first column. The initial level of expected inflation is set to its value at the end of the sample period.

Table 7: Means and standard deviations of the posterior parameter distributions

	2003		2004		2005		2006		2007		2008		2009		2010	
	mean	**std.dev.**	**mean**	**std.dev.**	**mean**	**std.dev.**	**mean**	**std.dev.**	**mean**	**std.dev.**	**mean**	**std.dev.**	**mean**	**std.dev.**	**mean**	**std.dev.**
simple VAR																
η_1	0.1785	0.0746	0.1606	0.0738	0.1635	0.0726	0.1653	0.0711	0.1629	0.0700	0.1697	0.0691	0.1606	0.0805	0.1680	0.0795
γ_1	0.0425	0.0199	0.0367	0.0196	0.0375	0.0192	0.0380	0.0187	0.0372	0.0184	0.0395	0.0181	0.0386	0.0211	0.0400	0.0209
η_2	-0.1085	0.0749	-0.0975	0.0738	-0.1084	0.0727	-0.1129	0.0710	-0.1136	0.0701	-0.1213	0.0695	-0.1258	0.0794	-0.1463	0.0812
γ_2	0.9738	0.0199	0.9774	0.0196	0.9741	0.0192	0.9727	0.0187	0.9726	0.0184	0.9701	0.0182	0.9678	0.0208	0.9633	0.0213
ω_{12}	-0.9439	0.0131	-0.9402	0.0136	-0.9383	0.0136	-0.9377	0.0135	-0.9356	0.0135	-0.9367	0.0131	-0.9385	0.0124	-0.9253	0.0146
ω_1^2	0.0048	0.0008	0.0048	0.0008	0.0047	0.0008	0.0046	0.0007	0.0045	0.0007	0.0044	0.0007	0.0060	0.0009	0.0060	0.0009
ω_2^2	0.0048	0.0008	0.0048	0.0008	0.0047	0.0008	0.0046	0.0007	0.0045	0.0007	0.0045	0.0007	0.0059	0.0009	0.0062	0.0009
extended VAR																
μ_1	0.2258	0.0890	0.2531	0.0865	0.2730	0.0842	0.2773	0.0828	0.2667	0.0815	0.2591	0.0815	0.1542	0.0919	0.1120	0.0900
β_1	-1.0475	2.2328	-3.0054	2.1627	-3.6181	2.0194	-3.6914	1.9557	-3.5677	1.9603	-3.3696	1.9064	1.5809	1.9149	3.6140	1.9325
γ	0.0533	0.0213	0.0554	0.0208	0.0593	0.0203	0.0602	0.0199	0.0574	0.0195	0.0561	0.0195	0.0399	0.0225	0.0323	0.0220
μ_2	0.0007	0.0004	0.0005	0.0004	0.0006	0.0003	0.0005	0.0003	0.0005	0.0003	0.0005	0.0003	0.0005	0.0003	0.0002	0.0003
β_2	0.9167	0.0439	0.9303	0.0443	0.9274	0.0419	0.9314	0.0406	0.9279	0.0406	0.9294	0.0417	0.9282	0.0415	0.9548	0.0424
μ_3	-0.1248	0.0742	-0.1031	0.0730	-0.1144	0.0718	-0.1200	0.0702	-0.1153	0.0687	-0.1137	0.0687	-0.1423	0.0795	-0.1615	0.0809
β_3	0.9690	0.0198	0.9755	0.0194	0.9723	0.0190	0.9708	0.0185	0.9720	0.0181	0.9722	0.0180	0.9635	0.0209	0.9593	0.0213
ρ_{12}	-0.3338	0.1068	-0.3532	0.1022	-0.3639	0.0988	-0.3589	0.0964	-0.3577	0.0945	-0.2956	0.0966	0.0802	0.1025	0.0833	0.1006
ρ_{13}	-0.1958	0.1156	-0.2026	0.1117	-0.1923	0.1093	-0.1715	0.1069	-0.2012	0.1035	-0.2060	0.1009	0.2013	0.0990	0.2166	0.0963
ρ_{14}	-0.9373	0.0148	-0.9262	0.0169	-0.9308	0.0154	-0.9316	0.0149	-0.9299	0.0148	-0.9306	0.0144	-0.9331	0.0136	-0.9045	0.0186
ρ_{23}	0.5614	0.0823	0.4924	0.0888	0.4946	0.0859	0.4713	0.0863	0.4334	0.0877	0.4175	0.0872	0.5917	0.0673	0.5473	0.0711
ρ_{24}	0.2966	0.1094	0.2773	0.1076	0.2993	0.1034	0.2972	0.1009	0.2964	0.0987	0.2251	0.1003	-0.0916	0.1023	-0.1259	0.0998
ρ_{34}	0.1953	0.1156	0.2237	0.1109	0.2212	0.1079	0.2010	0.1058	0.2363	0.1021	0.2291	0.1000	-0.1483	0.1010	-0.1793	0.0979
σ_1^2	0.0049	0.0009	0.0048	0.0008	0.0047	0.0008	0.0046	0.0007	0.0045	0.0007	0.0044	0.0007	0.0061	0.0009	0.0060	0.0009
σ_2^2	0.0009	0.0002	0.0010	0.0002	0.0009	0.0002	0.0009	0.0001	0.0009	0.0001	0.0009	0.0001	0.0012	0.0038	0.0245	0.0036
σ_3^2	0.0108	0.0019	0.0112	0.0019	0.0112	0.0018	0.0120	0.0019	0.0133	0.0021	0.0138	0.0021	0.0251	0.9008	6.3346	0.9295
σ_4^2	0.0050	0.0009	0.0049	0.0008	0.0049	0.0008	0.0047	0.0008	0.0046	0.0007	0.0046	0.0007	0.0060	0.0060	0.0060	0.0060

This table displays the means and standard deviations of the posterior distributions for the parameters of the VAR models in Equations (15) and (16). The parameters σ_2^2 and σ_3^2 have been multiplied by a factor 1,000. Estimation results are provided for quarterly samples starting in 1985 and ending in the first quarter of the years 2003 up to 2010, as indicated in the first column.

Table 8: Optimal stock allocations (in %) for agnostic investors with different beliefs and different investment horizons

	2003	2004	2005	2006	2007	2008	2009	2010
				simple VAR				
5	32.2	49.6	57.5	56.7	52.2	87.3	81.5	56.7
7	25.4	44.1	56.9	59.5	54.9	85.4	72.7	54.1
10	19.8	36.0	54.4	62.2	59.0	86.1	63.4	48.8
				extended VAR				
5	18.6	32.6	36.4	41.8	25.9	62.7	48.1	7.6
7	11.3	22.3	27.9	35.0	22.1	50.8	42.4	5.6
10	6.5	11.8	18.0	26.0	17.7	37.8	35.5	3.8

This table displays the optimal stock allocations (in % of initial real-term wealth) for agnostic investors with different beliefs. We consider investment horizons equal to five, seven and ten years, as indicated in the first column. The stock allocations in the upper part of the table correspond to an agnostic investor who makes investment decisions on the basis of a two-dimensional VAR model for stock returns and dividend yields (thus ignoring expected and unexpected inflation); see Equation (16). The lower part of the table displays the stock holdings of an agnostic investor who uses a four-dimensional VAR model for stock returns, expected and unexpected inflation and dividend yields; see Equation (15). The allocations are based on quarterly samples starting in 1985 and ending in the first quarter of the years 2003 up to 2010. The initial levels of the dividend yield and expected inflation are set to their values at the end of the sample period.

Appendix A Optimal stock allocations with power utility

We consider a power utility investor with risk aversion parameter ϕ. At time t, he wants to determine optimal portfolio wealth shares λ_t and $1-\lambda_t$ to be invested in the stock and an inflation-linked bond, respectively. Throughout, we assume that the investor follows a k-period buy-and-hold strategy.

We first observe that maximizing $\mathbb{E}_t\big[u(W_{t+k})\big]$ is equivalent to maximizing $\log \mathbb{E}_t\big[u(W_{t+k})\big]$ for $\phi \leq 1$ and to minimizing $\log\big[-\mathbb{E}_t\big[u(W_{t+k})\big]\big]$ for $\phi \geq 1$. Without loss of generality we assume that $\phi < 1$. Assuming log-normality of k-period real-term wealth, we have

$$\log \mathbb{E}_t\big[u(W_{t+k})\big] = (1-\phi)E_t[w_{t+k}] + (1/2)(1-\phi)^2\mathbb{V}\mathrm{ar}_t[w_{t+k}] - \log(1-\phi), \quad \text{(A.1)}$$

where $w_{t+k} = \log(W_{t+k})$. Observe that $w_{t+k} = r_{p,t}(k) + w_t$, with $r_{p,t}(k)$ the k-period continuously compounded real portfolio return. We can rewrite Equation (A.1) as

$$\log \mathbb{E}_t\big[u(W_{t+k})\big] = (1-\phi)\mathbb{E}_t[r_{p,t}(k)] + (1-\phi)w_t + (1/2)(1-\phi)^2\mathbb{V}\mathrm{ar}_t[w_{t+k}] - \log(1-\phi). \quad \text{(A.2)}$$

Maximizing the expression in Equation (A.2) is equivalent to maximizing

$$\log \mathbb{E}_t\big[u(W_{t+k})\big] = \mathbb{E}_t[r_{p,t}(k)] + (1/2)(1-\phi)\mathbb{V}\mathrm{ar}_t[r_{p,t}(k)]. \quad \text{(A.3)}$$

Since

$$\log \mathbb{E}_t\big[\exp(r_{p,t}(k))\big] = \mathbb{E}_t[r_{p,t}(k)] + (1/2)\mathbb{V}\mathrm{ar}_t[r_{p,t}(k)], \quad \text{(A.4)}$$

we can rewrite Equation (A.3) as

$$\log \mathbb{E}_t\big[u(W_{t+k})\big] = \mathbb{E}_t[1 + R_{p,t}(k)] - (\phi/2)\mathbb{V}\mathrm{ar}_t[r_{p,t}(k)], \quad \text{(A.5)}$$

where $R_t(k)$ denotes the simple net k-period portfolio return. For $\phi = 1$ the investor maximizes the expected log real portfolio returns, since in this case Equation (A.5) boils down to

$$\log \mathbb{E}_t\big[u(W_{t+k})\big] = \mathbb{E}_t[r_{p,t}(k)]. \tag{A.6}$$

For $\phi \leq 1$ the investor opts for a riskier portfolio, since a higher portfolio variance corresponds to a higher simple gross return (provided that the mean of the continuously compounded returns remains the same). For $\phi \geq 1$ the investor faces a trade-off between the mean and the variance of the portfolio return and chooses a less risky portfolio. Hence, the conditional mean and variance of the portfolio returns are crucial ingredients of the power utility framework with log-normal terminal wealth. Following Campbell et al. (2003), we approximate the continuously compounded real portfolio return by

$$r_{p,t}(k) \approx \alpha_t r_t(k) + (1-\alpha_t) r_{f,t}(k) + (1/2)\alpha_t(1-\alpha_t)\mathbb{V}\text{ar}_t[r_t(k)]. \tag{A.7}$$

Here α_t is the share invested in the stock at time t. The above approximation becomes more accurate for smaller k and it is exact in continuous time according to Itô's lemma. Using the above approximation, the optimal share invested in stocks is given by

$$\alpha_t = \frac{\mathbb{E}_t[r_t(k)] - r_{f,t}(k) + (1/2)\mathbb{V}\text{ar}_t[r_t(k)]}{\phi \mathbb{V}\text{ar}_t[r_t(k)]}. \tag{A.8}$$

With positive expected excess returns and $\phi > 0$, the optimal weight is a decreasing function of the conditional variance and an increasing function of the expected real (excess) return.