

NBER WORKING PAPER SERIES

ARE STOCKS REALLY LESS VOLATILE IN THE LONG RUN?

Lubos Pastor
Robert F. Stambaugh

Working Paper 14757
<http://www.nber.org/papers/w14757>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2009

We are grateful for comments from John Campbell, Peter Christoffersen, Darrell Duffie, Gene Fama, Wayne Ferson, Michael Halling, Cam Harvey, Anthony Lynch, Matt Richardson, Jeremy Siegel, Georgios Skoulakis, Pietro Veronesi, Luis Viceira, Rob Vishny, an associate editor, two anonymous referees, workshop participants at CFA Institute of Chicago, Columbia University, Comenius University, Cornell University, Erasmus University Rotterdam, George Mason University, Louisiana State University, Princeton University, Stanford University, Tilburg University, University of Amsterdam, University of California at Berkeley, University of California Los Angeles, University of Chicago, University of Pennsylvania, University of Southern California, University of Warwick, Washington University, WU Wien, and Yale University, and participants in the following conferences: 2010 American Finance Association, 2009 Western Finance Association, 2009 European Finance Association, 2009 NBER Summer Institute, 2009 Gerzensee, 2010 Chicago Booth Management Conference, 2010 Q Group, and the 2009 Symposium on Quantitative Methods in Finance at the University of Texas at Austin. We also gratefully acknowledge research support from the Q Group, and we thank Hyun Paul Lee for helpful research assistance. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2009 by Lubos Pastor and Robert F. Stambaugh. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Are Stocks Really Less Volatile in the Long Run?

Lubos Pastor and Robert F. Stambaugh

NBER Working Paper No. 14757

February 2009, Revised March 2011

JEL No. G11,G12,G23

ABSTRACT

According to conventional wisdom, annualized volatility of stock returns is lower when computed over long horizons than over short horizons, due to mean reversion induced by return predictability. In contrast, we find that stocks are substantially more volatile over long horizons from an investor's perspective. This perspective recognizes that parameters are uncertain, even with two centuries of data, and that observable predictors imperfectly deliver the conditional expected return. Mean reversion contributes strongly to reducing long-horizon variance, but it is more than offset by various uncertainties faced by the investor, especially uncertainty about the expected return. The same uncertainties also make target-date funds undesirable to a class of investors who would otherwise find them appealing.

Lubos Pastor

University of Chicago

Booth School of Business

5807 South Woodlawn Ave

Chicago, IL 60637

and NBER

lubos.pastor@chicagobooth.edu

Robert F. Stambaugh

Finance Department

The Wharton School

University of Pennsylvania

Philadelphia, PA 19104-6367

and NBER

stambaugh@wharton.upenn.edu

1. Introduction

Conventional wisdom views stock returns as less volatile over longer investment horizons. This view seems consistent with various empirical estimates. For example, using two centuries of U.S. equity returns, Siegel (2008) reports that variances realized over investment horizons of several decades are substantially lower than short-horizon variances on a per-year basis. Such evidence pertains to unconditional variance, but a similar message is delivered by studies that condition variance on information useful in predicting returns. Campbell and Viceira (2002, 2005), for example, report estimates of conditional variances that decrease with the investment horizon.

We find that stocks are actually *more* volatile over long horizons from an investor’s perspective. Investors condition on available information but realize their knowledge is limited in two key respects. First, even after observing 206 years of data (1802–2007), investors do not know the values of the parameters of the return-generating process, especially the parameters related to the conditional expected return. Second, investors recognize that observable “predictors” used to forecast returns deliver only an imperfect proxy for the conditional expected return, whether or not the parameter values are known. When viewed from this perspective, the return variance per year at a 50-year horizon is at least 1.3 times higher than the variance at a 1-year horizon.

Our main object of interest is the *predictive* variance of $r_{T,T+k}$, the k -period return starting at time T . Predictive variance, denoted by $\text{Var}(r_{T,T+k}|D_T)$, conditions on D_T , the data available to investors at time T . From an investor’s perspective, predictive variance is the relevant variance—the one suitable for portfolio decisions. Readers might be more familiar with *true* variance, which conditions on ϕ , the parameters of the return-generating process. Investors realize they do not know ϕ , and predictive variance incorporates that parameter uncertainty by conditioning only on D_T . In contrast, true variance conditions on ϕ , whether or not it also conditions on D_T . The true unconditional variance, $\text{Var}(r_{T,T+k}|\phi)$, is estimated by the usual sample variance, as in Siegel (2008). The true conditional variance, $\text{Var}(r_{T,T+k}|\phi, D_T)$, is estimated by Campbell and Viceira (2002, 2005). True variance is the more common focus of statistical inference. For example, an extensive literature uses variance ratios and other statistics to test whether $(1/k)\text{Var}(r_{T,T+k}|\phi)$ is the same for every investment horizon k .¹ We focus on $(1/k)\text{Var}(r_{T,T+k}|D_T)$ instead. That is, we compare long- and short-horizon predictive variances, which matter to investors. Investors might well infer from the data that the true variance is lower at long horizons, while at the same time assessing the predictive variance to be higher at long horizons.

¹A partial list of such studies includes Fama and French (1988), Poterba and Summers (1988), Lo and MacKinlay (1988, 1989), Richardson and Stock (1989), Kim, Nelson, and Startz (1991), and Richardson (1993).

The distinction between predictive variance and true variance is readily seen in the simple case where an investor knows the true variance of returns but not the true expected return. Uncertainty about the expected return contributes to the investor’s overall uncertainty about what the upcoming realized returns will be. Predictive variance includes that uncertainty, while true variance excludes it. Expected return is notoriously hard to estimate. Uncertainty about current expected return and about how expected return will change in the future is the key element of our story. This uncertainty plays an increasingly important role as the investment horizon grows, as long as investors believe that expected return is “persistent,” i.e., likely to take similar values across adjacent periods.

Under the traditional random-walk assumption that returns are distributed independently and identically (i.i.d.) through time, true return variance per period is equal at all investment horizons. Explanations for lower true variance at long horizons commonly focus on “mean reversion,” whereby a negative shock to the current return is offset by positive shocks to future returns, and vice versa. Our conclusion that stocks are more volatile in the long run obtains despite the presence of mean reversion. We show that mean reversion is only one of five components of long-run predictive variance:

- (i) i.i.d. uncertainty
- (ii) mean reversion
- (iii) uncertainty about future expected returns
- (iv) uncertainty about current expected return
- (v) estimation risk.

Whereas the mean-reversion component is strongly negative, the other components are all positive, and their combined effect outweighs that of mean reversion.

Of the four components contributing positively, the one making the largest contribution at long horizons reflects uncertainty about future expected returns. This component (iii) is often neglected in discussions of how return predictability affects long-horizon return variance. Such discussions typically highlight mean reversion, but mean reversion—and predictability more generally—require variance in the conditional expected return, which we denote by μ_t . That variance makes the future values of μ_t uncertain, especially in the more distant future periods, thereby contributing to the overall uncertainty about future returns. The greater the degree of predictability, the larger is the variance of μ_t and thus the greater is the relative contribution of uncertainty about future expected returns to long-horizon predictive variance.

Three additional components also make significant positive contributions to long-horizon predictive variance. One is simply the variance attributable to unexpected returns. Under an i.i.d.

assumption for unexpected returns, this variance makes a constant contribution to variance per period at all investment horizons. At long horizons, this component (i), though quite important, is actually smaller in magnitude than both components (ii) and (iii) discussed above.

Another component of long-horizon predictive variance reflects uncertainty about the current μ_t . Components (i), (ii), and (iii) all condition on the current value of μ_t . Conditioning on the current expected return is standard in long-horizon variance calculations using a vector autoregression (VAR), such as Campbell (1991) and Campbell, Chan, and Viceira (2003). In reality, though, an investor does not observe μ_t . We assume the investor observes the histories of returns and a given set of return predictors. This information is capable of producing only an imperfect proxy for μ_t , which in general reflects additional information. Pástor and Stambaugh (2009) introduce a predictive system to deal with imperfect predictors, and we use that framework to assess long-horizon predictive variance and capture component (iv). When μ_t is persistent, uncertainty about the current μ_t contributes to uncertainty about μ_t in multiple future periods, on top of the uncertainty about future μ_t 's discussed earlier.

The fifth and last component adding to long-horizon predictive variance, also positively, is one we label “estimation risk,” following common usage of that term. This component reflects the fact that, after observing the available data, an investor remains uncertain about the parameters of the joint process generating returns, expected returns, and the observed predictors. That parameter uncertainty adds to the overall variance of returns assessed by an investor. If the investor knew the parameter values, this estimation-risk component would be zero.

Parameter uncertainty also enters long-horizon predictive variance more pervasively. Unlike the fifth component, the first four components are non-zero even if the parameters are known to an investor. At the same time, those four components can be affected significantly by parameter uncertainty. Each component is an expectation of a function of the parameters, with the expectation evaluated over the distribution characterizing an investor's parameter uncertainty. We find that Bayesian posterior distributions of these functions are often skewed, so that less likely parameter values exert a significant influence on the posterior means, and thus on long-horizon predictive variance.

The effects of parameter uncertainty on the predictive variance of long-horizon returns are analyzed in previous studies, such as Stambaugh (1999), Barberis (2000), and Hoevenaars et al (2007). Barberis discusses how parameter uncertainty essentially compounds across periods and exerts stronger effects at long horizons. The above studies find that predictive variance is substantially higher than estimates of true variance that ignore parameter uncertainty. However, all three studies also find that long-horizon predictive variance is lower than short-horizon variance for the

horizons considered—up to 10 years in Barberis (2000), up to 20 years in Stambaugh (1999), and up to 50 years in Hoevenaars et al (2007).² In contrast, we often find that predictive variance even at a 10-year horizon is higher than at a 1-year horizon.

A key difference between our analysis and the above studies is our inclusion of uncertainty about the current expected return μ_t . The above studies employ VAR approaches in which observed predictors perfectly capture μ_t , whereas we consider predictors to be imperfect, as explained earlier. We compare predictive variances under perfect versus imperfect predictors, and find that long-run variance is substantially higher when predictors are imperfect. Predictor imperfection increases long-run variance both directly and indirectly. The direct effect, component (iv) of predictive variance, is large enough at a 10-year horizon that subtracting it from predictive variance leaves the remaining portion lower than the 1-year variance.

The indirect effect of predictor imperfection is even larger. It stems from the fact that predictor imperfection and parameter uncertainty interact—once predictor imperfection is admitted, parameter uncertainty is more important in general. This result occurs despite the use of informative prior beliefs about parameter values, as opposed to the non-informative priors used in the above studies. When μ_t is not observed, learning about its persistence and predictive ability is more difficult than when μ_t is assumed to be given by observed predictors. The effects of parameter uncertainty pervade all components of long-horizon returns, as noted earlier. The greater parameter uncertainty accompanying predictor imperfection further widens the gap between our analysis and the previous studies.³

Predictor imperfection can be viewed as omitting an unobserved predictor from the set of observable predictors used in a standard predictive regression. The degree of predictor imperfection can be characterized by the increase in the R-squared of that predictive regression if the omitted predictor were included. Even if investors assign a low probability to this increase being larger than 2% for annual returns, such modest predictor imperfection nevertheless exerts a substantial effect on long-horizon variance. At a 30-year horizon, for example, the predictive variance is 1.2 times higher than when the predictors are assumed to be perfect.

Our empirical results indicate that stocks should be viewed by investors as more volatile at

²Instead of actually reporting predictive variance, Barberis reports a closely related quantity: the asset allocation for a buy-and-hold, power-utility investor. His allocations for the 10-year horizon exceed those for short horizons, even when parameter uncertainty is incorporated.

³Schotman, Tschernig, and Budek (2008) find that if the predictors are fractionally integrated, long-horizon variance of stock returns can exceed short-horizon variance. With stationary predictors, though, they find long-horizon variance is smaller than short-horizon variance. By incorporating predictor imperfection as well as parameter uncertainty, we find that long-horizon variance exceeds short-horizon variance even when predictors are stationary.

long horizons. Corporate Chief Financial Officers (CFO's) indeed tend to exhibit such a view, as we discover by analyzing survey evidence reported by Ben-David, Graham, and Harvey (2010). In quarterly surveys conducted over eight years, Ben-David et al. ask CFO's to express confidence intervals for the stock market's annualized return over the next year and the next ten years. From the reported results of these surveys, we infer that the typical CFO views the annualized variance of ten-year returns to be at least twice the one-year variance.

The long-run volatility of stocks is of substantial interest to investors. Evidence of lower long-horizon variance is cited in support of higher equity allocations for long-run investors (e.g, Siegel, 2008) as well as the increasingly popular target-date mutual funds (e.g., Gordon and Stockton, 2006, Greer, 2004, and Viceira, 2008). These funds gradually reduce an investor's stock allocation by following a predetermined "glide path" that depends only on the time remaining until the investor's target date, typically retirement. When the parameters and conditional expected return are assumed to be known, we find that the typical glide path of a target-date fund closely resembles the pattern of allocations desired by risk-averse investors with utility for wealth at the target date. Once uncertainty about the parameters and conditional expected return is recognized, however, the same investors find the typical glide path significantly less appealing. Investors with sufficiently long horizons instead prefer glide paths whose initial as well as final stock allocations are substantially lower than those of investors with shorter horizons.

The remainder of the paper proceeds as follows. Section 2 derives expressions for the five components of long-horizon variance discussed above and analyzes their theoretical properties. Section 3 describes our empirical framework, which uses up to 206 years of data to implement two predictive systems that allow us to analyze various properties of long-horizon variance. Section 4 explores the five components of long-horizon variance using a predictive system in which the conditional expected return follows a first-order autoregression. Section 5 then gauges the importance of predictor imperfection using an alternative predictive system that includes an unobservable predictor. Section 6 discusses the robustness of our results. Section 7 returns to the above discussion of the distinction between an investor's problem and inference about true variance. Section 8 considers the implications of the CFO surveys reported by Ben-David et al. (2010). Section 9 analyzes investment implications of our results in the context of target-date funds. Section 10 summarizes our conclusions.

2. Long-horizon variance and parameter uncertainty

Let r_{t+1} denote the continuously compounded return from time t to time $t + 1$. We can write

$$r_{t+1} = \mu_t + u_{t+1}, \quad (1)$$

where μ_t denotes the expected return conditional on all information at time t and u_{t+1} has zero mean. Also define the k -period return from period $T + 1$ through period $T + k$,

$$r_{T,T+k} = r_{T+1} + r_{T+2} + \dots + r_{T+k}. \quad (2)$$

An investor assessing the variance of $r_{T,T+k}$ uses D_T , a subset of all information at time T . In our empirical analysis in Section 4, D_T consists of the full histories of returns as well as predictors that investors use in forecasting returns.⁴ Importantly, D_T typically reveals neither the value of μ_T in equation (1) nor the values of the parameters governing the joint dynamics of r_t , μ_t , and the predictors. Let ϕ denote the vector containing those parameter values.

This paper focuses on $\text{Var}(r_{T,T+k}|D_T)$, the predictive variance of $r_{T,T+k}$ given the investor's information set. Since the investor is uncertain about μ_T and ϕ , it is useful to decompose this variance as

$$\text{Var}(r_{T,T+k}|D_T) = \text{E}\{\text{Var}(r_{T,T+k}|\mu_T, \phi, D_T)|D_T\} + \text{Var}\{\text{E}(r_{T,T+k}|\mu_T, \phi, D_T)|D_T\}. \quad (3)$$

The first term in this decomposition is the expectation of the conditional variance of k -period returns. This conditional variance, which has been estimated by Campbell and Viceira (2002, 2005), is of interest only to investors who know the true values of μ_T and ϕ . Investors who do not know μ_T and ϕ are interested in the expected value of this conditional variance, and they also account for the variance of the conditional expected k -period return, the second term in equation (3). As a result, they perceive returns to be more volatile and, as we show below, they perceive disproportionately more volatility at long horizons. Whereas the conditional per-period variance of stock returns appears to decrease with the investment horizon, we show that $(1/k)\text{Var}(r_{T,T+k}|D_T)$, which accounts for uncertainty about μ_T and ϕ , increases with the investment horizon.

The potential importance of parameter uncertainty for long-run variance is readily seen in the special case where returns are i.i.d. with known variance σ^2 and unknown mean μ . In this case, the mean and variance of k -period returns conditional on μ are both linear in k : the mean is $k\mu$ and the variance is $k\sigma^2$. An investor who knows μ faces the same per-period variance, σ^2 , regardless

⁴We are endowing the investor with the same information set as the set that we use in our empirical analysis. In that sense, we are putting investors and econometricians on an equal footing, in the spirit of Hansen (2007).

of k . However, an investor who does not know μ faces more variance, and this variance increases with k . To see this, apply the variance decomposition from equation (3):

$$\begin{aligned}\text{Var}(r_{T,T+k}|D_T) &= \text{E}\{k\sigma^2|D_T\} + \text{Var}\{k\mu|D_T\} \\ &= k\sigma^2 + k^2\text{Var}\{\mu|D_T\},\end{aligned}\tag{4}$$

so that $(1/k)\text{Var}(r_{T,T+k}|D_T)$ increases with k . In fact, $(1/k)\text{Var}(r_{T,T+k}|D_T) \rightarrow \infty$ as $k \rightarrow \infty$. That is, an investor who believes that stock prices follow a random walk but who is uncertain about the unconditional mean μ views stocks as more volatile in the long run.

To assess the likely magnitude of this effect, consider the following back-of-the-envelope calculation. If uncertainty about μ is given by the standard error of the sample average return computed over T periods, or σ/\sqrt{T} , then $(1/k)\text{Var}(r_{T,T+k}|D_T) = \sigma^2(1 + k/T)$. With $k = 50$ years and $T = 206$ years, as in the sample that we use in Section 4, $(1 + k/T) = 1.243$, so the per-period predictive variance exceeds σ^2 by a quarter. Of course, if the sample mean estimate of μ is computed from a sample shorter than 206 years (e.g., due to concerns about nonstationarity), then uncertainty about μ is larger and the effect on predictive variance is even stronger.

When returns are predictable, so that μ_t is time-varying, $\text{Var}(r_{T,T+k}|D_T)$ can be above or below its value in the i.i.d. case. Predictability can induce mean reversion, which reduces long-run variance, but predictability also introduces uncertainty about additional quantities, such as future values of μ_t and the parameters that govern its behavior. It is not clear a priori whether predictability makes returns more or less volatile at long horizons, compared to the i.i.d. case. At sufficiently long horizons, uncertainty about the unconditional expected return will still dominate and drive $(1/k)\text{Var}(r_{T,T+k}|D_T)$ to infinity. At long horizons of relevance to investors, whether or not that per-period variance is higher than at short horizons is an empirical question that we explore.

In the rest of this section, we assume for simplicity that μ_t follows an AR(1) process,⁵

$$\mu_{t+1} = (1 - \beta)E_r + \beta\mu_t + w_{t+1}, \quad 0 < \beta < 1.\tag{5}$$

The AR(1) assumption for μ_t allows us to further decompose both terms on the right-hand side of equation (3), providing additional insights into the components of $\text{Var}(r_{T,T+k}|D_T)$. The AR(1) assumption also allows a simple characterization of mean reversion. Time variation in μ_t induces mean reversion in returns if the unexpected return u_{t+1} is negatively correlated with future values of μ_t . Under the AR(1) assumption, mean reversion requires a negative correlation between u_{t+1} and

⁵Our stationary AR(1) process for μ_t nests a popular model in which the stock price is the sum of a random walk and a positively autocorrelated stationary AR(1) component (e.g., Summers, 1986, Fama and French, 1988). In that special case, ρ_{uw} as well as return autocorrelations at all lags are negative. See the Appendix.

w_{t+1} , or $\rho_{uw} < 0$. If fluctuations in μ_t are persistent, then a negative shock in u_{t+1} is accompanied by offsetting positive shifts in the μ_{t+i} 's for multiple future periods, resulting in a stronger negative contribution to the variance of long-horizon returns.

2.1. Conditional variance

This section analyzes the conditional variance $\text{Var}(r_{T,T+k}|\mu_T, \phi, D_T)$, which is an important building block in computing the variance in equation (3). The conditional variance reflects neither parameter uncertainty nor uncertainty about the current expected return, since it conditions on both ϕ and μ_T . The parameter vector ϕ includes all parameters in equations (1) and (5): $\phi = (\beta, E_r, \rho_{uw}, \sigma_u, \sigma_w)$, where σ_u and σ_w are conditional standard deviations of u_{t+1} and w_{t+1} , respectively. Assuming that equations (1) and (5) hold and that the conditional covariance matrix of $[u_{t+1} \ w_{t+1}]$ is constant, $\text{Var}(r_{T,T+k}|\mu_T, \phi, D_T) = \text{Var}(r_{T,T+k}|\mu_T, \phi)$. Furthermore, we show in the Appendix that

$$\text{Var}(r_{T,T+k}|\mu_T, \phi) = k\sigma_u^2 \left[1 + 2\bar{d}\rho_{uw}A(k) + \bar{d}^2B(k) \right], \quad (6)$$

where

$$A(k) = 1 + \frac{1}{k} \left(-1 - \beta \frac{1 - \beta^{k-1}}{1 - \beta} \right) \quad (7)$$

$$B(k) = 1 + \frac{1}{k} \left(-1 - 2\beta \frac{1 - \beta^{k-1}}{1 - \beta} + \beta^2 \frac{1 - \beta^{2(k-1)}}{1 - \beta^2} \right) \quad (8)$$

$$\bar{d} = \left[\frac{1 + \beta}{1 - \beta} \frac{R^2}{1 - R^2} \right]^{1/2}, \quad (9)$$

and R^2 is the ratio of the variance of μ_t to the variance of r_{t+1} , based on equation (1).

The conditional variance in (6) consists of three terms. The first term, $k\sigma_u^2$, captures the well-known feature of i.i.d. returns—the variance of k -period returns increases linearly with k . The second term, containing $A(k)$, reflects mean reversion in returns arising from the likely negative correlation between realized returns and expected future returns ($\rho_{uw} < 0$), and it contributes negatively to long-horizon variance. The third term, containing $B(k)$, reflects the uncertainty about future values of μ_t , and it contributes positively to long-horizon variance. When returns are unpredictable, only the first term is present (because $R^2 = 0$ implies $\bar{d} = 0$, so the terms involving $A(k)$ and $B(k)$ are zero). Now suppose that returns are predictable, so that $R^2 > 0$ and $\bar{d} > 0$. When $k = 1$, the first term is still the only one present, because $A(1) = B(1) = 0$. As k increases, though, the terms involving $A(k)$ and $B(k)$ become increasingly important, because both $A(k)$ and $B(k)$ increase monotonically from 0 to 1 as k goes from 1 to infinity.

Figure 1 plots the variance in (6) on a per-period basis (i.e., divided by k), as a function of the investment horizon k . Also shown are the terms containing $A(k)$ and $B(k)$. It can be verified that $A(k)$ converges to 1 faster than $B(k)$. (See Appendix.) As a result, the conditional variance in Figure 1 is U-shaped: as k increases, mean reversion exerts a stronger effect initially, but uncertainty about future expected returns dominates eventually.⁶ The contribution of the mean reversion term, and thus the extent of the U-shape, is stronger when ρ_{uw} takes larger negative values. The contributions of mean reversion and uncertainty about future μ_{T+i} 's both become stronger as predictability increases. These effects are illustrated in Figure 2, which plots the same quantities as Figure 1, but for three different R^2 values. Note that a higher R^2 implies not only stronger mean reversion but also a more volatile μ_t , which in turn implies more uncertainty about future μ_{T+i} 's.

The key insight arising from Figures 1 and 2 is that, although mean reversion significantly reduces long-horizon variance, that reduction can be more than offset by uncertainty about future expected returns. Both effects become stronger as R^2 increases, but uncertainty about future expected returns prevails when R^2 is high. A high R^2 implies high volatility in μ_t and therefore high uncertainty about μ_{T+j} . In that case, long-horizon variance exceeds short-horizon variance on a per-period basis, even though ϕ and the current μ_T are assumed to be known. Uncertainty about ϕ and the current μ_T exerts a greater effect at longer horizons, further increasing the long-horizon variance relative to the short-horizon variance.

2.2. Components of long-horizon variance

The variance of interest, $\text{Var}(r_{T,T+k}|D_T)$, consists of two terms on the right-hand side of equation (3). The first term is the expectation of the conditional variance in equation (6), so each of the three terms in (6) is replaced by its expectation with respect to ϕ . (We need not take the expectation with respect to μ_T , since μ_T does not appear on the right in (6).) The interpretations of these terms are the same as before, except that now each term also reflects parameter uncertainty.

The second term on the right-hand side of equation (3) is the variance of the true conditional expected return. This variance is taken with respect to ϕ and μ_T . It can be decomposed into two components: one reflecting uncertainty about the current μ_T , or predictor imperfection, and the other reflecting uncertainty about ϕ , or “estimation risk.” (See the Appendix.) Let b_T and q_T

⁶Campbell and Viceira (2002, pp. 95–96) also model expected return as an AR(1) process, but they conclude that variance per period cannot increase with k when $\rho_{uw} < 0$. They appear to equate conditional variances of single-period returns across future periods, which would omit the uncertainty about future expected return.

denote the conditional mean and variance of the unobservable expected return μ_T :

$$b_T = \mathbb{E}(\mu_T | \phi, D_T) \quad (10)$$

$$q_T = \text{Var}(\mu_T | \phi, D_T). \quad (11)$$

The right-hand side of equation (3) can then be expressed as the sum of five components:

$$\begin{aligned} \text{Var}(r_{T,T+k} | D_T) = & \underbrace{\mathbb{E}\{k\sigma_u^2 | D_T\}}_{\text{i.i.d. uncertainty}} + \underbrace{\mathbb{E}\{2k\sigma_u^2 \bar{d}\rho_{uw} A(k) | D_T\}}_{\text{mean reversion}} + \underbrace{\mathbb{E}\{k\sigma_u^2 \bar{d}^2 B(k) | D_T\}}_{\text{future } \mu_{T+i} \text{ uncertainty}} \\ & + \underbrace{\mathbb{E}\left\{\left(\frac{1-\beta^k}{1-\beta}\right)^2 q_T | D_T\right\}}_{\text{current } \mu_T \text{ uncertainty}} + \underbrace{\text{Var}\left\{kE_r + \frac{1-\beta^k}{1-\beta}(b_T - E_r) | D_T\right\}}_{\text{estimation risk}}. \end{aligned} \quad (12)$$

Parameter uncertainty plays a role in all five components in equation (12). The first four components are expected values of quantities that are viewed as random due to uncertainty about ϕ , the parameters governing the joint dynamics of returns and predictors. (If the values of these parameters were known to the investor, the expectation operators could be removed from those four components.) Parameter uncertainty can exert a non-trivial effect on the first four components, in that the expectations can be influenced by parameter values that are unlikely but cannot be ruled out. The fifth component in equation (12) is the variance of a quantity whose randomness is also due to parameter uncertainty. In the absence of such uncertainty, the fifth component is zero, which is why we assign it the interpretation of estimation risk.

The estimation risk term includes the variance of kE_r , where E_r denotes the unconditional mean return. This variance equals $k^2 \text{Var}(E_r | D_T)$, so the per-period variance $(1/k) \text{Var}(r_{T,T+k} | D_T)$ increases at rate k . Similar to the i.i.d. case, if E_r is unknown, then the per-period variance grows without bounds as the horizon k goes to infinity. For finite horizons that are typically of interest to investors, however, the fifth component in equation (12) can nevertheless be smaller in magnitude than the other four components. In general, the k -period variance ratio, defined as

$$V(k) = \frac{(1/k) \text{Var}(r_{T,T+k} | D_T)}{\text{Var}(r_{T+1} | D_T)}, \quad (13)$$

can exhibit a variety of patterns as k increases. Whether or not $V(k) > 1$ at various horizons k is an empirical question.

3. Empirical framework: Predictive systems

It is commonly assumed that the conditional expected return μ_t is given by a linear combination of a set of observable predictors, x_t , so that $\mu_t = a + b'x_t$. This assumption is useful in many applications, but we relax it here because it understates the uncertainty faced by an investor assessing the variance of future returns. Any given set of predictors x_t is likely to be *imperfect*, in that μ_t is unlikely to be captured by any linear combination of x_t ($\mu_t \neq a + b'x_t$). The true expected return μ_t generally reflects more information than what we assume to be observed by the investor—the histories of r_t and x_t . To incorporate the likely presence of predictor imperfection, we employ a predictive system, defined in Pástor and Stambaugh (2009) as a state-space model in which r_t , x_t , and μ_t follow a VAR with coefficients restricted so that μ_t is the mean of r_{t+1} .⁷ As noted by Pástor and Stambaugh, a predictive system can also be represented as a VAR for r_t , x_t , and an unobserved additional predictor. We employ both versions here, as each is best suited to different dimensions of our investigation. Our two predictive systems are specified as follows:

System 1

$$r_{t+1} = \mu_t + u_{t+1} \quad (14)$$

$$x_{t+1} = \theta + Ax_t + v_{t+1} \quad (15)$$

$$\mu_{t+1} = (1 - \beta)E_r + \beta\mu_t + w_{t+1}. \quad (16)$$

System 2

$$r_{t+1} = a + b'x_t + \pi_t + u_{t+1} \quad (17)$$

$$x_{t+1} = \theta + Ax_t + v_{t+1} \quad (18)$$

$$\pi_{t+1} = \delta\pi_t + \eta_{t+1}. \quad (19)$$

In System 1, the conditional expected return μ_t is unobservable, and we assume $0 < \beta < 1$. System 2 includes π_t as an unobserved additional predictor of return, and we assume $0 < \delta < 1$. In both systems, the eigenvalues of A are assumed to lie inside the unit circle, and the vector containing the residuals of the three equations is assumed to be normally distributed, independently and identically across t .

System 1 is well suited for analyzing the components of predictive variance discussed in the previous section, because the AR(1) specification for μ_{t+1} in equation (16) is the same as that in equation (5). Pástor and Stambaugh (2009) provide a detailed analysis of System 1, and we apply

⁷State-space models have been used in a number of studies analyzing return predictability, including Conrad and Kaul (1988), Lamoureux and Zhou (1996), Johannes, Polson, and Stroud (2002), Ang and Piazzesi (2003), Brandt and Kang (2004), Dangl and Halling (2006), Duffee (2006), and Rytchkov (2007).

their econometric methodology in this study. In the next section, we investigate empirically the components of predictive variance using System 1.

System 2 is well suited for exploring the role of predictor imperfection in determining predictive variance. To see this, let σ_π^2 denote the variance of π_{t+1} in equation (19). As $\sigma_\pi^2 \rightarrow 0$, the predictors approach perfection, and equation (17) approaches the standard predictive regression,

$$r_{t+1} = a + b'x_t + e_{t+1}. \quad (20)$$

By examining results under various prior beliefs about the possible magnitudes of σ_π^2 , we can assess the effect of predictor imperfection on predictive variance. We do so in Section 5.

We conduct analyses using both annual and quarterly data. Our annual data consist of observations for the 206-year period from 1802 through 2007, as compiled by Siegel (1992, 2008). The return r_t is the annual real log return on the U.S. equity market, and x_t contains three predictors: the dividend yield on U.S. equity, the first difference in the long-term high-grade bond yield, and the difference between the long-term bond yield and the short-term interest rate.⁸ We refer to these quantities as the “dividend yield,” the “bond yield,” and the “term spread,” respectively. These three predictors seem reasonable choices given the various predictors used in previous studies and the information available in Siegel’s dataset. Dividend yield and the term spread have long been entertained as return predictors (e.g., Fama and French, 1989). Using post-war quarterly data, Pástor and Stambaugh (2009) find that the long-term bond yield, relative to its recent levels, exhibits significant predictive ability in predictive regressions. That evidence motivates our choice of the bond-yield variable used here. All three predictors exhibit significant predictive abilities in a predictive regression as in (20), with an R^2 in that regression of 5.6%.⁹ Our quarterly data consist of observations for the 220-quarter period from 1952Q1 through 2006Q4. We use the same three predictors in x_t as Pástor and Stambaugh (2009): dividend yield, CAY, and bond yield.¹⁰

4. Components of predictive variance (System 1)

This section uses the first predictive system, specified in equations (14) through (16), to empirically assess long-horizon return variance from an investor’s perspective. In the first subsection, we

⁸We are grateful to Jeremy Siegel for supplying these data. The long-term bond yield series is constructed from the yields of federal bonds and high-grade municipal bonds, as described in Siegel (1992).

⁹Details of the predictive regression results and the bootstrap significance tests are provided in an Internet Appendix available on the author’s websites.

¹⁰See that study for more detailed descriptions of the predictors. Our quarterly sample ends in 2006Q4 because the 2007 data on CAY of Lettau and Ludvigson (2001) are not yet available as of this writing. Our quarterly sample begins in 1952Q1, after the 1951 Treasury-Fed accord that made possible the independent conduct of monetary policy.

specify prior distributions for the system’s parameters and analyze the resulting posteriors. Those posterior distributions characterize the parameter uncertainty faced by an investor who conditions on essentially the entire history of U.S. equity returns. That uncertainty is incorporated in the Bayesian predictive variance, which is the focus of the second subsection. We analyze the five components of predictive variance and their dependence on the investment horizon. For this analysis, we report results using annual data. Results based on quarterly data are summarized later in Section 6; detailed results are reported in the Internet Appendix.

4.1. Priors and posteriors

For each of the three key parameters that affect multiperiod variance— ρ_{uw} , β , and R^2 —we implement the Bayesian empirical framework under three different prior distributions, displayed in Figure 3. The priors are assumed to be independent across parameters and follow the same functional forms as in Pástor and Stambaugh (2009). For each parameter, we specify a “benchmark” prior as well as two priors that depart from the benchmark in opposite directions but seem at least somewhat plausible as alternative specifications. When we depart from the benchmark prior for one of the parameters, we hold the priors for the other two parameters at their benchmarks, obtaining a total of seven different specifications of the joint prior for ρ_{uw} , β , and R^2 . We estimate the predictive system under each specification to explore the extent to which a Bayesian investor’s assessment of long-horizon variance is sensitive to prior beliefs.

The benchmark prior for ρ_{uw} , the correlation between expected and unexpected returns, has 97% of its mass below 0. This prior follows the reasoning of Pástor and Stambaugh (2009), who suggest that, a priori, the correlation between unexpected return and the innovation in expected return is likely to be negative. The more informative prior concentrates toward larger negative values, whereas the less informative prior essentially spreads evenly over the range from -1 to 1. The benchmark prior for β , the first-order autocorrelation in the annual expected return μ_t , has a median of 0.83 and assigns a low (2%) probability to β values less than 0.4. The two alternative priors then assign higher probability to either more persistence or less persistence. The benchmark prior for R^2 , the fraction of variance in annual returns explained by μ_t , has 63% of its mass below 0.1 and relatively little (17%) above 0.2. The alternative priors are then either more concentrated or less concentrated on low values. These priors on the true R^2 are shown in Panel C of Figure 3. Panel D displays the corresponding implied priors on the “observed” R^2 —the fraction of variance in annual real returns explained by the predictors. Each of the three priors in Panel D is implied by those in Panel C, while holding the priors for ρ_{uw} and β at their benchmarks and specifying non-informative priors for the degree of imperfection in the predictors. Observe that the benchmark

prior for the observed R^2 has much of its mass below 0.05.

We compute posterior distributions for the parameters using the Markov Chain Monte Carlo (MCMC) method discussed in Pástor and Stambaugh (2009). These posteriors summarize the parameter uncertainty faced by an investor after updating the priors using the 206-year history of equity returns and predictors. Figure 4 plots the posteriors corresponding to the priors plotted in Figure 3. The posteriors of β , shown in Panel B of Figure 4, reveal substantial persistence in the conditional expected return μ_t . The posterior modes are about 0.9, regardless of the prior, and β values smaller than 0.7 seem very unlikely. Comparing the posteriors with the priors in Figure 3, we see that the data shift the prior beliefs in the direction of higher persistence. The posteriors of the true R^2 , displayed in Panel C, lie to the right of the corresponding priors. For example, for the benchmark prior, the prior mode for the true R^2 is less than 0.05, while the posterior mode is nearly 0.1. The data thus shift the priors in the direction of greater predictability. The same message is conveyed by the posteriors of the observed R^2 , plotted in Panel D.

The posteriors of ρ_{uw} are displayed in Panel A of Figure 4. These posteriors are more concentrated toward larger negative values than any of the three priors of ρ_{uw} , suggesting strong mean reversion in the data. The posteriors are similar across the three priors, consistent with observed autocorrelations of annual real returns and the posteriors of R^2 and β discussed above. Equations (1) and (5) imply that the autocovariances of returns are given by

$$\text{Cov}(r_t, r_{t-k}) = \beta^{k-1} (\beta \sigma_\mu^2 + \sigma_{uw}), \quad k = 1, 2, \dots, \quad (21)$$

where $\sigma_\mu^2 = \sigma_w^2 / (1 - \beta^2)$. From (21) we can also obtain the autocorrelations of returns,

$$\text{Corr}(r_t, r_{t-k}) = \beta^{k-1} \left(\beta R^2 + \rho_{uw} \sqrt{(1 - R^2) R^2 (1 - \beta^2)} \right), \quad k = 1, 2, \dots, \quad (22)$$

by noting that $\sigma_\mu^2 = R^2 \sigma_r^2$ and that $\sigma_u^2 = (1 - R^2) \sigma_r^2$. The posterior modes of ρ_{uw} in Figure 4 are about -0.9, and the posterior modes of R^2 and β are about 0.1 and 0.9, as observed earlier. Evaluating (22) at those values gives autocorrelations starting at -0.028 for $k = 1$ and then increasing gradually toward 0 as k increases. Such values are statistically indistinguishable from the observed autocorrelations of annual real returns in our sample.¹¹

¹¹The first five autocorrelations in our 206-year sample are 0.02, -0.17, -0.04, 0.01, and -0.10. To assess the compatibility of these sample autocorrelations with our predictive system, we proceed as follows. We first draw the full set of system parameters from their posterior distribution. Using these parameters, we simulate a 206-year sample of returns by drawing the error terms in equations (14) and (16) from their joint normal distribution. We then compute the first five autocorrelations for this simulated sample. Repeating this procedure for many posterior draws of parameters, we obtain many sets of sample autocorrelations simulated from the predictive system. These simulated sets form a five-dimensional probability density because there are five autocorrelations. We then consider a five-dimensional grid of autocorrelation values, spaced 0.03 apart, splitting the parameter space into a finite number

Panel A of Figure 5 plots the prior and posterior distributions for the R^2 in a regression of the conditional mean μ_t on the three predictors in x_t . This R^2 quantifies the degree of imperfection in the predictors ($R^2 = 1$ if and only if the predictors are perfect), which plays a key role in our analysis. Both distributions are obtained under the benchmark prior from Figure 3. The prior distribution for R^2 is rather noninformative, assigning nontrivial probability mass to the whole $(0, 1)$ interval. In contrast, the posterior distribution is substantially tighter, indicating relevant information in the data. This posterior reveals a substantial degree of predictor imperfection, in that the density's mode is about 0.3, and values above 0.8 have near-zero probability.

Further perspective on the predictive abilities of the individual predictors is provided by Panel B of Figure 5. This panel plots the posteriors of the partial correlations between μ_t and each predictor, obtained under the benchmark priors.¹² Dividend yield exhibits the strongest relation to expected return, with the posterior for its partial correlation ranging between 0 and 0.9 and having a mode around 0.6. Most of the posterior mass for the term spread's partial correlation lies above zero, but there is little posterior mass above 0.5. The bond yield's marginal contribution is the weakest, with much of the posterior density lying between -0.2 and 0.2. In the multiple regression of returns on the three predictors, described at the end of Section 3., all predictors (rescaled to have unit variances) have comparable OLS slope coefficients and t-statistics. When compared to those estimates, the posteriors in Panel B indicate that dividend yield is more attractive as a predictor but that bond yield is less attractive. These differences are consistent with the predictors' autocorrelations and the fact that the posterior distribution of β , the autocorrelation of μ_t , centers around 0.9. The autocorrelations for the three predictors are 0.92 for dividend yield, 0.65 for the term spread, and -0.04 for the bond yield. The bond yield's low autocorrelation makes it look less correlated with μ_t , whereas dividend yield's higher autocorrelation makes it look more like μ_t .

4.2. Multiperiod predictive variance and its components

Each of the five components of multiperiod return variance in equation (12) is a moment of a quantity evaluated with respect to the distribution of the parameters ϕ , conditional on the information

of five-dimensional 'buckets'. We calculate the empirical frequency F with which the bucket containing the observed set of autocorrelations (0.02, -0.17, -0.04, 0.01, -0.10) obtains in our simulations. Finally, we compute the p -value as the fraction of the simulated sets of autocorrelations that fall in buckets whose empirical frequency is smaller than F . The p -value based on 300,000 simulations is 37%, indicating that the predictive system cannot be rejected based on sample autocorrelations.

¹²The partial correlation of a predictive variable with μ_t is informative about the variable's predictive power for returns, but it does not necessarily measure the variable's importance for portfolio decisions. For a rebalancing investor, the contemporaneous correlation of the variable with stock return is important for determining the hedging demand for the stock. We would like to thank one of the referees for this valid observation.

D_T available to an investor at time T . In our Bayesian empirical setting, D_T consists of the 206-year history of returns and predictors, and the distribution of parameters is the posterior density given that sample. Draws of ϕ from this density are obtained via the MCMC procedure and then used to evaluate the required moments of each of the components in equation (12). The sum of those components, $\text{Var}(r_{T,T+k}|D_T)$, is the Bayesian predictive variance of $r_{T,T+k}$.

Figure 6 displays the predictive variance and its five components for horizons of $k = 1$ through $k = 50$ years, computed under the benchmark priors. The values are stated on a per-year basis (i.e., divided by k). The predictive variance (Panel A) increases significantly with the investment horizon, with the per-year variance exceeding the one-year variance by about 45% at a 30-year horizon and about 80% at a 50-year horizon. This is the main result of the paper.

The five variance components, displayed in Panel B of Figure 6, reveal the sources of the greater predictive variance at long horizons. Over a one-year horizon ($k = 1$), virtually all of the variance is due to the i.i.d. uncertainty in returns, with uncertainty about the current μ_T and parameter uncertainty also making small contributions. Mean reversion and uncertainty about future μ_t 's make no contribution for $k = 1$, but they become quite important for larger k . Mean reversion contributes negatively at all horizons, consistent with $\rho_{uw} < 0$ in the posterior (cf. Figure 4), and the magnitude of this contribution increases with the horizon. Nearly offsetting the negative mean reversion component is the positive component due to uncertainty about future μ_t 's. At longer horizons, the magnitudes of both components exceed the i.i.d. component, which is flat across horizons. At a 10-year horizon, the mean reversion component is nearly equal in magnitude to the i.i.d. component. At a 30-year horizon, both mean reversion and future- μ_t uncertainty are substantially larger in magnitude than the i.i.d. component. In fact, the mean reversion component is larger in magnitude than the overall predictive variance.

Both estimation risk and uncertainty about the current μ_T make stronger positive contributions to predictive variance as the investment horizon lengthens. At the 30-year horizon, the contribution of estimation risk is about two thirds of the contribution of the i.i.d. component. Uncertainty about the current μ_T , arising from predictor imperfection, makes the smallest contribution among the five components at long horizons, but it still accounts for almost a quarter of the total predictive variance at the 30-year horizon.

Table 1 reports the predictive variance at horizons of 25 and 50 years under various prior distributions for ρ_{uw} , β , and R^2 . For each of the three parameters, the prior for that parameter is specified as one of the three alternatives displayed in Figure 3, while the prior distributions for the other two parameters are maintained at their benchmarks. Also reported in Table 1 is the ratio of the long-horizon predictive variance to the one-year variance, as well as the contribution of each

of the five components to the long-horizon predictive variance.

Across the different priors in Table 1, the 25-year variance ratio ranges from 1.15 to 1.42, and the 50-year variance ratio ranges from 1.45 to 1.96. The variance ratios exhibit the greatest sensitivity to prior beliefs about R^2 . The “loose” prior beliefs that assign higher probability to larger R^2 values produce the lowest variance ratios. When returns are more predictable, mean reversion makes a stronger negative contribution to variance, but uncertainty about future μ_t 's makes a stronger positive contribution. Those two components are the largest in absolute magnitude. The next largest is the positive contribution from i.i.d. uncertainty, which declines as the prior on R^2 moves from tight to loose. Recall that i.i.d. uncertainty is the posterior mean of $k\sigma_u^2$. This posterior mean declines as the prior on R^2 loosens up because greater posterior density on high values of R^2 necessitates less density on high values of $\sigma_u^2 = (1 - R^2)\sigma_r^2$, given that the sample is informative about the unconditional return variance σ_r^2 . Prior beliefs about ρ_{uw} and β have a smaller effect on the predictive variance and its components.¹³

In sum, when viewed by an investor whose prior beliefs lie within the wide range of priors considered here, stocks are considerably more volatile at longer horizons. The greater volatility obtains despite the presence of a large negative contribution from mean reversion.

5. Perfect predictors versus imperfect predictors (System 2)

This section uses the second predictive system, given in equations (17) through (19), to investigate the extent to which long-run variance is affected by predictor imperfection. Recall that predictor imperfection in System 2 is equivalent to $\sigma_\pi^2 > 0$. Incorporating predictor imperfection is a key difference between our analysis and the studies by Stambaugh (1999) and Barberis (2000), which analyze the effects of parameter uncertainty on long-run equity volatility. Those studies model expected return as $\mu_t = a + b'x_t$, so that the observed predictors deliver expected return perfectly if the parameters a and b are known. The latter “perfect-predictor” assumption yields the predictive regression in (20), which obtains as the limit in System 2 when σ_π^2 approaches zero. Combining the predictive regression in (20) with the VAR for x_t in (18) then delivers implications for long-run variance in the perfect-predictor setting, as in Stambaugh (1999) and Barberis (2000).

To assess the importance of predictor imperfection, we compute predictive variances under

¹³This relative insensitivity to prior beliefs about ρ_{uw} and β appears to be specific to the long sample of real equity returns. Greater sensitivity to prior beliefs appears if returns in excess of the short-term interest rate are used instead, or if quarterly returns on a shorter and more recent sample period are used. In all of these alternative samples, we obtain variance results that lead to the same qualitative conclusions.

various informative prior beliefs about σ_π . Non-informative prior beliefs are specified for all other parameters of the predictive system except δ , the autocorrelation of the additional unobserved predictor.¹⁴ When using the annual data, we specify the prior distribution for δ to be the same as the benchmark prior in System 1 for β , the autocorrelation of the conditional mean. We shift the prior for δ somewhat closer to 1.0 when using the quarterly data, since a given persistence for the expected annual return is likely to correspond to a higher persistence at the quarterly frequency.¹⁵

We specify three different priors for σ_π . One of the priors has all of its mass at $\sigma_\pi = 0$, which is equivalent to an assumption of perfect predictors. The remaining two priors are displayed in the uppermost panels of Figure 7. Panel A shows the priors used with annual data, and Panel B shows those for the quarterly data. The latter densities are shifted closer to zero, consistent with the higher frequency. Updating these priors with the data produces the corresponding posterior densities for σ_π shown in Panels C and D. The posteriors for σ_π shift noticeably to the left versus the priors, indicating that the sample information plays a nontrivial role in resolving some of the uncertainty about predictor imperfection.

Investors' posterior beliefs about predictor imperfection can also be characterized in terms of ΔR^2 , defined as the “true” R^2 for predicting one-period returns—the R^2 when conditioning on both x_t and π_t —minus the “observed” R^2 when conditioning only on x_t . Panels E and F of Figure 7 show the posteriors for ΔR^2 . From these plots we see that, after updating with the sample data, investors in our setting believe predictor imperfection to be rather modest. For example, the specification with less predictor imperfection (solid line) has the bulk of the posterior mass below $\Delta R^2 = 0.02$ for annual data. In other words, after seeing the data, an investor in that case believes it is fairly unlikely that an unobserved predictor could raise the R^2 by more than two percent. With quarterly data, the corresponding posterior for ΔR^2 concentrates on even smaller values.

Even when investors assess potential predictor imperfection to be relatively modest, the imperfection has important consequences for the predictive variance of long-horizon returns. Predictive variances for horizons up to 50 years are shown in Panel A of Figure 8 for the annual data, while Panel B shows the corresponding results for the quarterly data. The importance of recognizing predictor imperfection emerges clearly from these results. In Panel A, the predictive variances at the longest horizons are about 1.3 times higher when predictor imperfection is recognized than when predictors are assumed to be perfect. For the quarterly results in Panel B, that ratio is well over 2.0.

¹⁴The Internet Appendix provides details of the Bayesian procedures, including the specification of priors and the calculation of predictive variances.

¹⁵With the annual data, the prior for δ is a truncated normal, where the mean and standard deviation of the non-truncated distribution are 0.99 and 0.25. The latter values are 0.99 and 0.15 with the quarterly data.

We also see in Figure 8 that predictive variances are substantially greater at long horizons than at short horizons, once predictor imperfection is recognized. Thus, the results for System 2 deliver the same overall message as the earlier results for System 1. In Panel A, using annual data, the predictive variance at the 50-year horizon is 1.4–1.5 times the 1-year variance, depending on the degree of predictor imperfection. In Panel B, using quarterly data, the 50-year variance is 1.3–1.4 times the 1-year variance.

Stambaugh (1999) and Barberis (2000) investigate the effects of parameter uncertainty using data beginning in 1952, the same year that our quarterly data begin. With these data, predictor imperfection plays an especially large role—more than doubling the variance at long horizons. With perfect predictors, consistent with Stambaugh and Barberis, predictive variance is substantially lower at long horizons: the 50-year variance ratio is then 0.6. In contrast, when predictor imperfection is incorporated, the 50-year variance ratio is 1.3–1.4, as observed above. Thus, when using post-1951 data, accounting for predictor imperfection rather dramatically reverses the answer to the question of whether stocks are less volatile in the long run.

We also see that the findings of Stambaugh and Barberis, indicating stocks are less volatile at longer horizons even after incorporating parameter uncertainty, do not obtain over the longer 206-year period. The predictive variances in Panel E are actually higher at long horizons, given perfect predictors, with a 50-year variance ratio just below 1.2. In all of our results, however, admitting predictor imperfection produces long-run variance that substantially exceeds not only short-run variance but also long-run variance computed assuming perfect predictors.

6. Robustness

6.1. Alternative samples

Our main empirical message—that long-run predictive variance of stock returns exceeds short-run variance—is robust to various sample specifications for both predictive systems.¹⁶ First, we extend the results for System 1 to the quarterly data included in the results for System 2. We adjust the prior distributions in System 1 to reflect the different data frequency, shifting the priors for R^2 and ρ_{uw} to the left and for β to the right.¹⁷ We find that the results with the quarterly data are

¹⁶Detailed results are reported in the Internet Appendix.

¹⁷We shift the prior on R^2 to the left because return predictability is likely to be weaker in quarterly data than in annual data. It is well known that in the presence of persistence in the conditional expected return, there is more predictability at lower data frequencies. We also shift the prior on ρ_{uw} to the left because the correlation between expected and unexpected returns is likely to be less negative at lower frequencies. Given stationarity in expected

even stronger than those with our annual data. Using the benchmark priors, the 25-year predictive variance is 92% larger than the 1-year variance, and the 50-year predictive variance is nearly 3 times the 1-year variance.

Second, instead of using real returns, we compute excess stock returns by subtracting the short-term interest rate from the realized stock return, and we then repeat the analyses for both predictive systems using both annual and quarterly data. The results are similar to those with real returns: all of the 50-year predictive variances exceed short-run variance by substantial amounts. Third, instead of using three predictors, we use only one, dividend yield, and repeat the analyses for both predictive systems using both annual and quarterly data. The results are again similar to the original three-predictor results: consistently higher predictive variances at long horizons.

Fourth, we conduct subperiod analyses for the results based on annual data. For both predictive systems, we split the 1802–2007 sample in half and estimate the predictive variances separately as of the ends of both subperiods. Under the same priors used in Figures 6 and 8, the predictive variance per period rises monotonically with the horizon under both systems in the first subperiod. In the second subperiod, the predictive variance rises monotonically under System 2, while under System 1 it exhibits a U-shape with respect to the horizon. In the latter case, the variance decreases through a horizon of 7 years but thereafter increases, exceeding the 1-year variance beyond an 18-year horizon. That is, the negative effect of mean reversion prevails at short horizons, but the combined positive effects of estimation risk and uncertainty about current and future μ_t 's prevail at long horizons. For both subperiods and both predictive systems, long-horizon predictive variance exceeds short-run variance across all specifications: the 50-year variance ratio is at least 1.25 under System 1 and at least 1.8 under System 2.

6.2. Model uncertainty

In general, investors are uncertain about whether expected return is a linear function of a set of observed predictors. In our setting of predictor imperfection, that uncertainty admits the possibility that an unobserved predictor also plays a role. Another dimension of uncertainty about expected return is whether one or more *observed* predictors is necessary. As a simple case, consider an investor who rules out an unobserved predictor but is uncertain about which observed predictors belong in the predictive regression that delivers expected return. The latter case of “model uncertainty” is analyzed by Avramov (2002) and Cremers (2002). We analyze predictive variance in this

returns, stock returns measured over increasingly long periods are likely to be increasingly driven by cash flow news as opposed to discount rate news. Finally, we shift the prior on β to the right because a given persistence in the expected annual return is likely to correspond to a higher persistence at the quarterly frequency.

setting in order to consider an alternative dimension of uncertainty about expected return.

As explained earlier, our data include three observed predictors in both the annual and quarterly samples. Therefore, for each sample, there are eight (2^3) possible models that represent different subsets of the three predictors (including the case of no predictors, i.e. constant expected return). To this set of models we apply the model-uncertainty framework of Avramov (2002). We change from his data to ours but otherwise follow his methodology and specifications, which include the assignment of equal prior probabilities across the possible models.¹⁸ Figure 9 reports the resulting predictive variances, which incorporate not only uncertainty about the parameters within each model but also uncertainty about which of the eight models best captures expected return. We see that predictive variance per period increases with the investment horizon for both the annual and quarterly data. For the annual data, the annualized variance at a 50-year horizon is about 1.24 times the one-year variance; for the quarterly data, the predictive variance per-quarter at a 200-quarter horizon is about 1.75 times the one-quarter variance. We thus see that our study's main conclusion—higher predictive variance per period for longer horizons—also obtains from this alternative perspective on uncertainty about expected return.

6.3. Time-varying volatility

Our implementation of predictive systems assumes that the covariance matrix of the disturbances is constant over time. This assumption may seem unappealing, given evidence of time-varying volatility reported in a large literature on that topic. The assumption offers two advantages for this study. First, it permits a more tractable framework for exploring the importance of parameter uncertainty and predictor imperfection for long-horizon volatility. We show that much of long-horizon volatility is induced by various aspects of uncertainty about expected returns, such as uncertainty about the current and future values of μ_t as well as about the parameters characterizing the process for μ_t . Uncertainty related to μ_t affects the perception of returns over many future periods; as a result, this uncertainty exerts an increasingly large effect on multiperiod volatility as the investment horizon increases. It is well known that μ_t is difficult to estimate, and this difficulty is highlighted once we recognize that predictors are imperfect. All of these arguments would remain valid if we allowed the covariance matrix of the disturbances to vary over time.

The second advantage of the constant-covariance-matrix assumption is that it allows us to abstract from fluctuations in short-run volatility that would complicate the question of whether stocks

¹⁸We refer the reader to Avramov (2002) for details, including the procedure for calculating predictive variance. He does not report variances but instead reports initial buy-and-hold asset allocations for size/book-to-market portfolios for horizons up to ten years.

are more volatile in the long run. To see the latter point, consider a period (such as the fall of 2008) when the current short-run volatility greatly exceeds its typical level. When looking forward from that point in time, investors almost surely see stocks as less volatile over longer investment horizons, due to the well-documented mean reversion in short-run volatility. Conversely, when short-run volatility is unusually low, investors may view stocks as more volatile in the long run simply because they expect volatility to increase toward its long-run mean. Such observations seem less interesting than asking whether stocks are less volatile over long horizons, abstracting from effects that can flip the answer back and forth through time. This question is also the focus of previous studies, cited earlier, that address long-horizon versus short-horizon equity volatility.

Allowing time-varying volatility need not change the analytical results in Section 2. To see this, suppose there is time variation in the conditional covariance matrix of $\kappa_t = [u_t \ v'_t \ w_t]$, the vector of residuals in System 1. Let Σ_t denote the conditional covariance matrix at time t of κ_{t+1} . It seems plausible to assume that, if $\Sigma_t = \Sigma$ at a given time t , then

$$E_t(\kappa_{t+i}\kappa'_{t+i}) = \Sigma \quad \text{for all } i > 0. \quad (23)$$

Such a property is satisfied, for example, by a stationary first-order multivariate GARCH process of the form

$$\text{vech}(\Sigma_t) = c_0 + C_1 \text{vech}(\kappa_t \kappa'_t) + C_2 \text{vech}(\Sigma_{t-1}), \quad (24)$$

where $\text{vech}(\cdot)$ stacks the columns of the lower triangular part of its argument. With (23), the conditional variance of the k -period return in equation (6) is unchanged, provided we interpret it as $\text{Var}(r_{T,T+k} | \mu_T, \phi, \Sigma_T = \Sigma)$. The introduction of parameter uncertainty is also unchanged, under the interpretation that Σ is uncertain but that, whatever it is, it also equals Σ_T . Setting $\Sigma_T = \Sigma$ removes horizon effects due to the mean reversion in Σ_T discussed earlier. If Σ_T were instead low relative to Σ , for example, then the reversion of future Σ_{T+i} s to Σ could also contribute to long-run volatility. Setting $\Sigma_T = \Sigma$ excludes such a contribution, producing a cleaner assessment of long-run volatility.

Time variation in volatility could potentially matter for long-horizon investing by inducing hedging demands. In a setting with dynamic rebalancing, investors could find it valuable to adjust their stock allocations for the purpose of hedging against adverse movements in volatility. Chacko and Viceira (2005) estimate the magnitude of the volatility-induced hedging demands by calibrating a model in which the inverse of volatility follows a simple mean-reverting process. They find that hedging demands are very small, due to insufficient variability and persistence in volatility. In reaching their conclusion, Chacko and Viceira assume that their parameter estimates are equal

to the parameters' true values. If parameter uncertainty were taken into account, the volatility-induced hedging demands could potentially be larger. We do not analyze hedging demands since our portfolio analysis in Section 9 considers a predetermined asset allocation policy. Nonetheless, we view the analysis of volatility-induced hedging demands in the presence of parameter uncertainty as an interesting topic for future research.

7. Predictive variance versus true variance

This section provides further perspective on our results by distinguishing between two different measures of variance: predictive variance and true variance. The “predictive variance,” our main object of interest thus far, is the variance from the perspective of an investor who conditions on the historical data but remains uncertain about the true values of the parameters. The “true variance” is defined as the variance conditional on the true parameter values. The predictive variance and the true variance coincide if the data history is infinitely long, in which case the parameters are estimated with infinite precision. Estimates of the true variance can be relevant in some applications, such as option pricing, but the predictive variance is relevant for portfolio decisions.

When conducting inference about the true variance, a commonly employed statistic is the sample long-horizon variance ratio. Values of such ratios are often less than 1 for stocks, suggesting lower unconditional variances per period at long horizons. Figure 10 plots sample variance ratios for horizons of 2 to 50 years computed with the 206-year sample of annual real log stock returns analyzed above. The calculations use overlapping returns and unbiased variance estimates.¹⁹ Also plotted are percentiles of the variance ratio's Monte Carlo sampling distribution under the null hypothesis that returns are i.i.d. normal. That distribution exhibits positive skewness and has nearly 60% of its mass below 1. The realized value of 0.28 at the 30-year horizon attains a Monte Carlo p -value of 0.01, supporting the inference that the true 30-year variance ratio lies below 1 (setting aside the multiple-comparison issues of selecting one horizon from many). Panel A of Figure 11 plots the posterior distribution of the 30-year ratio for true unconditional variance, based on the benchmark priors and System 1. Even though the posterior mean of this ratio is 1.34, the distribution is positively skewed and 63% of the posterior probability mass lies below one. We thus see that the variance ratio statistic in a frequentist setting and the posterior distribution in a Bayesian setting both favor the inference that the true unconditional variance ratio is below 1.

Inference about true unconditional variance ratios is of limited relevance to investors, for two

¹⁹Each ratio is computed as $\overline{VR}(q)$ in equation (2.4.37) of Campbell, Lo, and MacKinlay (1997).

reasons. First, even if the parameters and the conditional mean μ_T were known, the unconditional variance would not be the appropriate measure from an investor's perspective, because conditional variance is more relevant when returns are predictable. The ratio of true unconditional variances can be less than 1 while the ratio of true conditional variances exceeds 1, or vice versa. At a horizon of $k = 30$ years, for example, parameter values of $\beta = 0.60$, $R^2 = 0.30$, and $\rho_{uw} = -0.55$ imply a ratio of 0.90 for unconditional variances but 1.20 for conditional variances.²⁰

The second and larger point is that inference about true variance, conditional or unconditional, is distinct from assessing the predictive variance perceived by an investor who does not know the parameters. This distinction can be drawn clearly in the context of the variance decomposition,

$$\text{Var}(r_{T,T+k}|D_T) = \text{E}\{\text{Var}(r_{T,T+k}|\phi, D_T)|D_T\} + \text{Var}\{\text{E}(r_{T,T+k}|\phi, D_T)|D_T\}. \quad (25)$$

The variance on the left-hand side of (25) is the predictive variance. The quantity inside the expectation in the first term, $\text{Var}(r_{T,T+k}|\phi, D_T)$, is the true conditional variance, relevant only to an investor who knows the true parameter vector ϕ (but not μ_T , thus maintaining predictor imperfection). The data can imply that this *true* variance is probably *lower* at long horizons than at short horizons while also implying that the *predictive* variance is *higher* at long horizons. In other words, investors who observe D_T can infer that if they were told the true parameter values, they would probably assess 30-year variance to be less than 1-year variance. These investors realize, however, that they do not know the true parameters. As a consequence, they evaluate the posterior mean of the true conditional variance, the first term in (25). That posterior mean can exceed the most likely values of the true conditional variance, because the posterior distribution of the true variance can be skewed (we return to this point below). Moreover, investors must add to that posterior mean the posterior variance of the true conditional mean, the second term in (25), which is the same as the estimation-risk term in equation (12). In a sense, investors do conduct inference about true variance—they compute its posterior mean—but they realize that estimate is only part of predictive variance.

The results based on our 206-year sample illustrate how predictive variance can be higher at long horizons while true variance is inferred to be most likely higher at short horizons. Panel B of Figure 11 plots the posterior distribution of the variance ratio

$$V^*(k) = \frac{(1/k)\text{Var}(r_{T,T+k}|\phi, D_T)}{\text{Var}(r_{T+1}|\phi, D_T)}, \quad (26)$$

for $k = 30$ years. The posterior probability that this ratio of true variances lies below 1 is 76%, and

²⁰The relation between the ratios of conditional and unconditional variances is derived in the Appendix. Campbell and Viceira (2002, p. 96) state that the unconditional variance ratio is always greater than the conditional ratio, but it appears they equate single-period conditional and unconditional variances in reaching that conclusion.

the posterior mode is below 0.5. In contrast, recall that 30-year predictive variance is substantially greater than 1-year variance, as shown earlier in Figure 6 and Table 1.

The true conditional variance $\text{Var}(r_{T,T+k}|\phi, D_T)$ is the sum of four quantities, the first four components in equation (12) with the expectations operators removed. The posterior distributions of those quantities (not shown to save space) exhibit significant asymmetries. As a result, less likely values of these quantities exert a disproportionate effect on the posterior means and, therefore, on the first term of the predictive variance in (25). The components reflecting uncertainty about current and future μ_t are positively skewed, so their contributions to predictive variance exceed what they would be if evaluated at the most likely parameter values. This feature of parameter uncertainty also helps drive predictive variance above the most likely value of true variance.

8. Long-horizon variance: Survey evidence

Our empirical results show investors should view stocks as more volatile over long horizons than over short horizons. Corporate CFO's indeed appear to exhibit such a view, as can be inferred from survey results reported by Ben-David, Graham, and Harvey (2010). Their survey asks each CFO to give the 10th and 90th percentiles of a confidence interval for the annualized (average) excess equity return to be realized over the upcoming 10-year period. The same question is asked for a 1-year horizon. For each horizon (k), the authors use the 10th and 90th percentiles to approximate $\text{Var}(\bar{r}_k)$, the variance of the CFO's perceived distribution of the annualized return. The resulting standard deviations are then averaged across CFO's. If we treat the averaged standard deviations as those perceived by a "typical" CFO, we can infer the typical CFO's views about long-horizon variance.

The relation between $\text{Var}(\bar{r}_k)$ and the annualized variance of the k -year return, $(1/k)\text{Var}(r_{T,T+k})$, which is our object of interest, must obey

$$\begin{aligned} (1/k)\text{Var}(r_{T,T+k}) &= (1/k)\text{Var}\left(\sum_{i=1}^K r_{T+i}\right) \\ &= (1/k)\text{Var}(k\bar{r}_k) \\ &= k\text{Var}(\bar{r}_k). \end{aligned} \tag{27}$$

If CFO's perceive stocks as equally volatile at all horizons, as in the standard i.i.d. setting with no parameter uncertainty, then $(1/k)\text{Var}(r_{T,T+k}) = \text{Var}(r_{T,T+1})$ and $\text{Var}(\bar{r}_k) = \text{Var}(r_{T,T+1})/k$. In that case, the perceived standard deviation of the 1-year return should be 3.2 ($=\sqrt{10}$) times the perceived standard deviation of the annualized 10-year return. In the survey results reported by

Ben-David et al., we observe that the ratios of 1-year standard deviation to the 10-year standard deviation are substantially below 3.2. Across 33 quarterly surveys from the first quarter of 2002 through the first quarter of 2010, the ratio ranges from 1.25 to 2.14, and its average value is 1.54. Even the maximum ratio of 2.14 implies

$$\frac{\text{Var}(\bar{r}_1)}{\text{Var}(\bar{r}_{10})} = (2.14)^2, \quad (28)$$

or, applying (27), a 10-year variance ratio given by

$$\frac{(1/10)\text{Var}(r_{T,T+10})}{\text{Var}(r_{T,T+1})} = \frac{10}{(2.14)^2} = 2.18, \quad (29)$$

as compared to the value of 1.0 when stocks are equally volatile over long and short horizons. In other words, the typical CFO appears to view stock returns as having at least twice the variance over a 10-year horizon than over a 1-year horizon.

9. Target-date funds

This section explores the long-run riskiness of stocks from the perspective of a very popular investment strategy. Target-date funds, also known as life-cycle funds, represent one of the fastest-growing segments of the investment industry. Since the inception of these funds in the mid-1990's, their assets have grown to about \$280 billion in 2010, including a net cash inflow of \$42 billion during the tumultuous year 2008. About 87% of target-date fund assets are held in retirement accounts as of third-quarter 2010 (Investment Company Institute, 2011).

Target-date funds follow a predetermined asset allocation policy that gradually reduces the stock allocation as the target date approaches, with the aim of providing a more conservative asset mix to investors approaching retirement.²¹ A predetermined allocation policy is not optimal because it sacrifices the ability to rebalance in response to future events, an ability analyzed in numerous studies of dynamic asset allocation.²² We venture off the well-trod path of that literature to consider a long-horizon strategy that, while suboptimal in theory, has become important in practice. We do not attempt to explain why so many real-world investors desire a predetermined path for their asset allocations. We simply take that fact as given and analyze the asset allocation problem within that setting. This focus also seems natural in the context of our study, since long-horizon equity volatility is relevant for investors making long-horizon equity decisions.

²¹See Viceira (2008) for a more detailed discussion of target-date funds.

²²Recent examples include Balduzzi and Lynch (1999), Barberis (2000), Brandt, Goyal, Santa-Clara, and Stroud (2005), Brandt, Santa-Clara, and Valkanov (2009), Detemple, Garcia, Rindesbacher (2003), Lynch and Balduzzi (2000), and Lynch (2001), among others. Wachter (2010) provides a review of the asset-allocation literature.

To analyze target-date funds using a simple model, we consider an investor who can invest in two assets, the stock market and a real riskless asset. The investor's horizon is K years, and his utility for end-of-horizon wealth W_K is given by $W_K^{1-A}/(1-A)$. The investor commits at the outset to a predetermined investment strategy in which the stock allocation evolves linearly from the first-period allocation w_1 to the final-period allocation w_K . The investor solves for the values of w_1 and w_K within the $(0, 1)$ interval to maximize expected utility. The investor assumes that the conditional expected stock return at the beginning of each horizon, μ_T , is equal to the unconditional expected return E_r , while treating E_r as uncertain. This specification removes the effect that a non-zero value of $\mu_T - E_r$ would have on the investor's desired pattern of stock allocations over the investment horizon. We solve the problem numerically, setting relative risk aversion A to 8 and the riskless real rate to 2% per year.

Target-date funds are often motivated by arguments related to human capital and labor income. A typical argument goes as follows.²³ Human capital is bond-like as it offers a steady stream of labor income. Younger people have more human capital because they stand to collect labor income over a longer time period. Younger people thus have a larger implicit position in bonds. To balance that position, younger people should invest a bigger fraction of their financial wealth in stocks, and they should gradually reduce their stock allocation as they grow older.

We consider two frameworks that differ in their treatment of labor income. In the first framework, presented in Section 9.1, the investor invests an initial nest egg and does not invest additional savings from any labor income. In the second framework, presented in Section 9.2, the investor also saves a fraction of his labor income. Both frameworks lead to the same conclusions regarding the effects of parameter uncertainty on the stock allocations of long-horizon investors.

9.1. No Savings from Labor Income

In this subsection, we assume that the investor derives no savings from any labor income. The investor simply begins with initial financial wealth W_0 , which subsequently evolves as follows:

$$W_{t+1} = W_t [1 + w_t r_{S,t+1} + (1 - w_t) r_f] , \quad (30)$$

where $r_{S,t}$ is the simple stock return in year t and r_f is the risk-free rate.²⁴

²³See, for example, Bodie, Merton, and Samuelson (1992), Viceira (2001), Cocco, Gomes, and Maenhout (2005), and Gordon and Stockton (2006). Other recent studies that analyze portfolio choice in the presence of labor income include Gomes and Michaelides (2005), Benzoni, Collin-Dufresne, and Goldstein (2007), Gomes, Kotlikoff, and Viceira (2008), and Lynch and Tan (2009), among others.

²⁴There is no adjustment for intermediate consumption since the investor is concerned only about terminal wealth.

Panels A and B of Figure 12 plot the investor’s optimal initial and final stock allocations, w_1 (solid line) and w_K (dashed line), for investment horizons ranging from 1 to 30 years. In Panel A, parameter uncertainty is ignored, in that the parameters characterizing the return process are treated as known and equal to their posterior means. In Panel B, parameter uncertainty is incorporated by using the posterior distributions. These come from our baseline setting: System 1 implemented on the 1802–2007 sample with three predictors and the benchmark prior.

The optimal allocations in Panel A of Figure 12 are strikingly similar to those selected by real-world target-date funds. The initial allocation w_1 decreases steadily as the investment horizon shortens, declining from about 85% at long horizons such as 25 or 30 years to about 30% at the one-year horizon, whereas the final allocation w_K is roughly constant at about 30-40% across all horizons. Investors in real-world target-date funds similarly commit to a stock allocation schedule, or “glide path,” that decreases steadily to a given level at the target date. The final stock allocation in a target-date fund does not depend on when investors enter the fund, but the initial allocation does—it is higher for investors entering longer before the target date. Not only the patterns but also the magnitudes of the optimal allocations in Panel A resemble those of target-date funds. For example, Viceira (2008) reports that the target-date funds offered by Fidelity and Vanguard reduce their stock allocations from 90% at long horizons to about 30% at short horizons. In addition, Vanguard’s stock allocations equal 90% for all horizons of 25 years or longer (see Viceira’s Figure 5.2), which corresponds nicely to the relatively flat portion of the solid line in Panel A.²⁵ In short, target-date funds seem appealing to investors who maximize expected power utility of wealth at the target date and who ignore parameter uncertainty.

In contrast, target-date funds do not appear desirable if the same investors incorporate parameter uncertainty, as shown in Panel B. For short investment horizons, the results look similar to those in Panel A, but for longer horizons, neither w_1 nor w_K are roughly invariant to the horizon; instead, they both decrease with K . For example, an investor with a 15-year horizon chooses to glide from $w_1 = 62\%$ to $w_{15} = 33\%$, but an investor with a 30-year horizon chooses lower stock allocations, gliding from $w_1 = 57\%$ to $w_{30} = 7\%$. The long-horizon stock allocations are lower in Panel B because investors perceive disproportionately more parameter uncertainty at long horizons.

²⁵Our simplification of target-date funds does not impose the constraint, common in practice, that all investors with the same time remaining in their horizons also have the same allocation.

9.2. Labor Income

In this subsection, we assume that the investor saves a positive fraction of his labor income each year. The investor's financial wealth evolves as follows:

$$W_{t+1} = W_t [1 + w_t r_{S,t+1} + (1 - w_t) r_f] + s L_{t+1}, \quad (31)$$

where L_t denotes labor income and s is the savings rate. We assume a constant savings rate, abstracting from the fact that investors may benefit from dynamically adjusting their savings rates over time. A constant savings rate is consistent with the fact that the predominant use of target-date funds is in employer-sponsored retirement plans, where both the employer and employee contributions are typically predetermined fractions of income. We set $s = 2.20\%$, which is the average annual ratio of aggregate personal saving to personal income over the past 5 years (2005–2009), as reported by the Bureau of Economic Analysis.

We assume the following simple process for labor income growth:

$$L_{t+1}/L_t - 1 = \xi (43 - \text{age}_t) + \epsilon_{t+1}, \quad (32)$$

where ξ is a constant, age_t denotes the investor's age in year t , and ϵ_{t+1} is drawn randomly from $N(0, \sigma_\epsilon^2)$. We set $\sigma_\epsilon = 0.08$, which is equal to the estimate of the annualized standard deviation of wage income growth reported by Heaton and Lucas (2000). The motivation for the age-related term in equation (32) is the evidence that expected labor income exhibits a hump-shaped pattern over a typical investor's lifecycle. For example, Figure 1 in Cocco, Gomes, and Maenhout (2005) shows that labor income is an inverse-U-shape function of age, for each of three different groups of households sorted by their education level. To capture the concave pattern in the level of labor income, we assume that the growth rate of labor income is a linearly decreasing function of age. We calibrate this function to the middle line in Cocco et al's Figure 1, according to which expected labor income grows until age 43 and declines thereafter. We set $\xi = 0.0043$, so that initial labor income growth at age 20 is 10%, as in Cocco et al's Figure 1. We assume that the investor retires at age 65, which is also the end of his investment horizon, so that $\text{age}_t = 65 - K + t$.

Note that labor income growth in equation (32) is uncorrelated with stock market returns. This assumption is motivated by the evidence that the correlation between wage growth and the stock market is generally close to zero. For example, Heaton and Lucas (2000) report a correlation of -0.07, and Cocco, Gomes, and Maenhout (2005) report correlations ranging from -0.02 to 0.01 across three different education levels. However, the assumption of zero correlation is not necessary for our conclusions. In an earlier version of the paper, we modeled labor income growth

as a convex combination of returns on the stock market and the T-bill, and we found that our conclusions were unaffected by relatively large changes in the weight on the stock market.

To capture the fact that younger people (those with higher values of K) tend to have less financial wealth, we specify the initial ratio of financial wealth to labor income, denoted by $F_K = W_0/L_0$, as a decreasing function of horizon K . Given the retirement age of 65, F_K is the ratio of financial wealth to labor income for an investor with $\text{age}_0 = 65 - K$. We specify F_K as

$$F_K = \exp\left(-\frac{4}{45}K\right). \quad (33)$$

The function in equation (33) is empirically motivated by data from the 2007 Panel Study of Income Dynamics (PSID) compiled by the University of Michigan. For all ages between 20 and 65, we compute the median ratio of financial wealth to labor income across all households headed by a person of that age.²⁶ The natural logarithm of this median ratio is an approximately linear function of age, and its value is about -4 for age 20 and about 0 for age 65. Adopting this linear approximation and recognizing that $K = 65 - \text{age}_0$, we quickly obtain equation (33).

Panels C and D of Figure 12 plot the investor's optimal initial and final stock allocations, w_1 and w_K , as a function of the investment horizon. These panels are constructed in the same way as Panels A and B, except that the investor's financial wealth follows equation (31) rather than equation (30). Parameter uncertainty is incorporated in Panel D but not in Panel C.

Similar to Panel A, the optimal allocations in Panel C look very much like those adopted by target-date funds. The initial allocation w_1 decreases from 100% at horizons longer than 15 years to about 30% at the one-year horizon, whereas the final allocation w_K is roughly constant at 30-40% across all horizons. Target-date funds thus seem appealing to investors who ignore parameter uncertainty even if those investors have labor income savings. In contrast, Panel D shows that target-date funds do not seem appealing if the same investors incorporate parameter uncertainty. For horizons longer than 23 years, both w_1 and w_K decrease with K . For example, an investor with a 23-year horizon chooses to glide from $w_1 = 100\%$ to $w_{23} = 14\%$, whereas an investor with a 30-year horizon glides from $w_1 = 93\%$ to $w_{30} = 3\%$. Echoing our earlier observation in the absence of labor income savings, the long-horizon stock allocations are lower in Panel D because investors perceive more parameter uncertainty at long horizons.

In Figure 12, investors always optimally choose downward-sloping glide paths, $w_K < w_1$, for all $K > 1$. This choice is not driven by mean reversion; $w_K < w_1$ remains optimal even if mean reversion is eliminated by setting $\rho_{ww} = 0$. Instead, the driving force is that future expected returns

²⁶The financial wealth of each household is computed by adding up items S805, S811, S815, and S819 in PSID.

μ_{T+j} are unknown and likely to be persistent. As j increases, the future values μ_{T+j} become increasingly uncertain from the perspective of investors at time T . As a result, the future returns $r_{T+j+1} = \mu_{T+j} + u_{T+j+1}$ become increasingly volatile from the investors' perspective. In other words, investors perceive distant future returns to be more volatile than near-term returns. Facing the need to predetermine their future allocations, investors commit to invest less in stocks in the more uncertain distant future. This simple logic shows that neither mean reversion nor human capital are necessary to justify downward-sloping glide paths. If investors must commit to a fixed schedule of future stock allocations, they will choose lower allocations at longer horizons simply because they view single-period stock returns as more volatile at longer horizons.

The results in Figure 12 demonstrate how parameter uncertainty makes target-date funds undesirable when they would otherwise be virtually optimal for investors who desire a predetermined asset-allocation policy. It would be premature, however, to conclude that parameter uncertainty makes target-date funds undesirable to such investors in all settings. The above analysis abstracts from many important considerations faced by investors, such as intermediate consumption, housing, etc. Our objective in this section is simply to illustrate how parameter uncertainty can reduce the stock allocations of long-horizon investors, consistent with our results about long-horizon volatility.

10. Conclusions

We use predictive systems and up to 206 years of data to compute long-horizon variance of real stock returns from the perspective of an investor who recognizes that parameters are uncertain and predictors are imperfect. Mean reversion reduces long-horizon variance considerably, but it is more than offset by other effects. As a result, long-horizon variance substantially exceeds short-horizon variance on a per-year basis. A major contributor to higher long-horizon variance is uncertainty about future expected returns, a component of variance that is inherent to return predictability, especially when expected return is persistent. Estimation risk is another important component of predictive variance that is higher at longer horizons. Uncertainty about current expected return, arising from predictor imperfection, also adds considerably to long-horizon variance. Accounting for predictor imperfection is key in reaching the conclusion that stocks are substantially more volatile in the long run. Overall, our results show that long-horizon stock investors face more volatility than short-horizon investors, in contrast to previous research.

In computing predictive variance, we assume that the parameters of the predictive system remain constant over 206 years. Such an assumption, while certainly strong, is motivated by our

objective to be conservative in treating parameter uncertainty. This uncertainty, which already contributes substantially to long-horizon variance, would generally be even greater under alternative scenarios in which investors would effectively have less information about the current values of the parameters. There is of course no guarantee that using a longer sample is conservative. In principle, for example, the predictability exhibited in a given shorter sample could be so much higher that both parameter uncertainty as well as long-run predictive variance would be lower. However, when we examine a particularly relevant shorter sample, a quarterly post-war sample spanning 55 years, we find that our main results get even stronger.

Changing the sample is only one of many robustness checks performed in the paper. We have considered a number of different prior distributions and modeling choices, reaching the same conclusion. Nonetheless, we cannot rule out the possibility that our conclusion would be reversed under other priors or modeling choices. In fact, we already know that if expected returns are modeled in a particularly simple way, assuming perfect predictors, then investors who rely on the post-war sample view stocks as less volatile in the long run. By continuity, stocks will also appear less volatile if only a very small degree of predictor imperfection is admitted a priori. Our point is that this traditional conclusion about long-run volatility is reversed in a number of settings that we view as more realistic, even when the degree of predictor imperfection is relatively modest.

Our finding that predictive variance of stock returns is higher at long horizons makes stocks less appealing to long-horizon investors than conventional wisdom would suggest. A clear illustration of such long-horizon effects emerges from our analysis of target-date funds. We demonstrate that a simple specification of the investment objective makes such funds appealing in the absence of parameter uncertainty but less appealing in the presence of that uncertainty. However, one must be cautious in drawing conclusions about the desirability of stocks for long-horizon investors in settings with additional risky assets, such as nominal bonds, additional life-cycle considerations, such as intermediate consumption, and optimal dynamic saving and investment decisions. Investigating asset-allocation decisions in such settings, while allowing the higher long-run stock volatility to enter the problem, is beyond the scope of this study but offers interesting directions for future research.

Appendix

A.1. Derivation of the conditional variance $\text{Var}(r_{T,T+k} | \mu_T, \phi)$

We can rewrite the AR(1) process for μ_t in equation (5) as an MA(∞) process

$$\mu_t = \mathbf{E}_r + \sum_{i=0}^{\infty} \beta^i w_{t-i}, \quad (\text{A1})$$

given our assumption that $0 < \beta < 1$. From (1) and (A1), the return k periods ahead is equal to

$$r_{T+k} = (1 - \beta^{k-1})\mathbf{E}_r + \beta^{k-1}\mu_T + \sum_{i=1}^{k-1} \beta^{k-1-i} w_{T+i} + u_{T+k}. \quad (\text{A2})$$

The multiperiod return from period $T + 1$ through period $T + k$ is then

$$r_{T,T+k} = \sum_{i=1}^k r_{T+i} = k\mathbf{E}_r + \frac{1 - \beta^k}{1 - \beta} (\mu_T - \mathbf{E}_r) + \sum_{i=1}^{k-1} \frac{1 - \beta^{k-i}}{1 - \beta} w_{T+i} + \sum_{i=1}^k u_{T+i}. \quad (\text{A3})$$

The conditional variance of the k -period return can be obtained from equation (A3) as

$$\begin{aligned} \text{Var}(r_{T,T+k} | \mu_T, \phi) &= k\sigma_u^2 + \frac{\sigma_w^2}{(1 - \beta)^2} \left[k - 1 - 2\beta \frac{1 - \beta^{k-1}}{1 - \beta} + \beta^2 \frac{1 - \beta^{2(k-1)}}{1 - \beta^2} \right] \\ &\quad + \frac{2\sigma_{uw}}{1 - \beta} \left[k - 1 - \beta \frac{1 - \beta^{k-1}}{1 - \beta} \right]. \end{aligned} \quad (\text{A4})$$

Equation (A4) can then be written as in equations (6) to (9), where \bar{d} arises from the relation

$$\sigma_w^2 = \sigma_\mu^2(1 - \beta^2) = \sigma_r^2 R^2(1 - \beta^2) = (\sigma_u^2 / (1 - R^2)) R^2(1 - \beta^2). \quad (\text{A5})$$

A.2. Properties of $A(k)$ and $B(k)$

1. $A(1) = 0, B(1) = 0$
2. $A(k) \rightarrow 1$ as $k \rightarrow \infty, B(k) \rightarrow 1$ as $k \rightarrow \infty$
3. $A(k+1) > A(k) \forall k, B(k+1) > B(k) \forall k$
4. $A(k) \geq B(k) \forall k$, with a strict inequality for all $k > 1$
5. $0 \leq A(k) < 1, 0 \leq B(k) < 1$
6. $A(k)$ converges to one more quickly than $B(k)$

Properties 1 and 2 are obvious. Properties 3 and 4 are proved below. Property 5 follows from Properties 1–3. Property 6 follows from Properties 1–4.

Proof that $A(k+1) > A(k) \forall k$:

$$\begin{aligned} A(k+1) &= 1 + \frac{1}{k+1} \left[-1 - \beta(1 + \beta + \dots + \beta^{k-2} + \beta^{k-1}) \right] \\ &= 1 + \frac{k}{k+1} \frac{1}{k} \left[-1 - \beta(1 + \beta + \dots + \beta^{k-2} + \beta^{k-1}) \right] \\ &= 1 + \frac{k}{k+1} \left[A(k) - 1 - \frac{\beta^k}{k} \right], \end{aligned}$$

which exceeds $A(k)$ if and only if $A(k) < 1 - \beta^k$. This is indeed true because

$$A(k) = 1 - \frac{1}{k} - \frac{1}{k} \left[\beta^1 + \dots + \beta^{k-1} \right] = 1 - \frac{1}{k} \left[\beta^0 + \beta^1 + \dots + \beta^{k-1} \right] < 1 - \frac{1}{k} \left[k\beta^k \right] = 1 - \beta^k.$$

Proof that $B(k+1) > B(k) \forall k$:

$$\begin{aligned} B(k+1) &= 1 + \frac{1}{k+1} \left[-1 - 2\beta(1 + \beta + \dots + \beta^{k-2} + \beta^{k-1}) + \beta^2(1 + \beta^2 + \dots + (\beta^2)^{k-2} + (\beta^2)^{k-1}) \right] \\ &= 1 + \frac{k}{k+1} \frac{1}{k} \left[\left\{ -1 - 2\beta(1 + \beta + \dots + \beta^{k-2}) + \beta^2(1 + \beta^2 + \dots + (\beta^2)^{k-2}) \right\} - 2\beta^k + \beta^{2k} \right] \\ &= 1 + \frac{k}{k+1} \left[B(k) - 1 + \frac{1}{k} \left(-2\beta^k + \beta^{2k} \right) \right], \end{aligned}$$

which exceeds $B(k)$ if and only if $B(k) < 1 + \beta^{2k} - 2\beta^k$. This is indeed true because

$$\begin{aligned} B(k) &= 1 - 2\frac{1}{k} + \frac{1}{k} - 2\frac{1}{k} \left(\beta + \dots + \beta^{k-2} + \beta^{k-1} \right) + \frac{1}{k} \left(\beta^2 + \dots + (\beta^2)^{k-2} + (\beta^2)^{k-1} \right) \\ &= 1 + \frac{1}{k} \left[\left((\beta^2)^0 - 2\beta^0 \right) + \left((\beta^2)^1 - 2\beta^1 \right) + \dots + \left((\beta^2)^{k-1} - 2\beta^{k-1} \right) \right] \\ &< 1 + \frac{1}{k} \left[k \left((\beta^2)^k - 2\beta^k \right) \right] \\ &= 1 + \beta^{2k} - 2\beta^k, \end{aligned}$$

where the inequality follows from the fact that the function $f(x) = (\beta^2)^x - 2\beta^x$ is increasing in x (because $f'(x) = 2(\ln\beta)\beta^x(\beta^x - 1) > 0$, for $0 < \beta < 1$).

Proof that $A(k) > B(k) \forall k > 1$:

$$\begin{aligned} B(k) - A(k) &= \frac{1}{k} \left[\beta^2 \frac{1 - \beta^{2(k-1)}}{1 - \beta^2} - \beta \frac{1 - \beta^{k-1}}{1 - \beta} \right] = \frac{1}{k} \left[\beta^2 + \dots + (\beta^2)^{k-1} - \left(\beta + \dots + \beta^{k-1} \right) \right] \\ &= \frac{1}{k} \sum_{i=1}^{k-1} \left(\beta^{2i} - \beta^i \right) = \frac{1}{k} \sum_{i=1}^{k-1} \beta^i \left(\beta^i - 1 \right) < 0. \end{aligned}$$

A.3. Decomposition of $\text{Var}\{E(r_{T,T+k}|\mu_T, \phi, D_T)|D_T\}$

Let $E_{T,k} = E(r_{T,T+k}|\mu_T, \phi, D_T)$. The variance of $E_{T,k}$ given D_T can be decomposed as

$$\text{Var}\{E_{T,k}|D_T\} = E\{\text{Var}[E_{T,k}|\phi, D_T]|D_T\} + \text{Var}\{E[E_{T,k}|\phi, D_T]|D_T\}. \quad (\text{A6})$$

To simplify each term on the right-hand side, observe from equations (1), (2), and (5) that

$$\begin{aligned} E_{T,k} &= E(r_{T+1} + r_{T+2} + \dots + r_{T+k}|\mu_T, \phi, D_T) \\ &= E(\mu_T + \mu_{T+1} + \dots + \mu_{T+k-1}|\mu_T, \phi) \\ &= kE_r + \frac{1 - \beta^k}{1 - \beta}(\mu_T - E_r). \end{aligned} \quad (\text{A7})$$

Taking the first and second moments of (A7), using (10) and (11), then gives

$$E[E_{T,k}|\phi, D_T] = kE_r + \frac{1 - \beta^k}{1 - \beta}(b_T - E_r) \quad (\text{A8})$$

$$\text{Var}[E_{T,k}|\phi, D_T] = \left(\frac{1 - \beta^k}{1 - \beta}\right)^2 q_T. \quad (\text{A9})$$

Substituting (A8) and (A9) into (A6) then gives the fourth and fifth terms in (12), using (3).

A.4. Relation between conditional and unconditional variance ratios

The unconditional variance (which does not condition on μ_T) is given by

$$\begin{aligned} \text{Var}(r_{T,T+k}|\phi) &= E[\text{Var}(r_{T,T+k}|\mu_T, \phi, D_T)|\phi] + \text{Var}[E(r_{T,T+k}|\mu_T, \phi, D_T)|\phi] \\ &= \text{Var}(r_{T,T+k}|\mu_T, \phi) + \left(\frac{1 - \beta^k}{1 - \beta}\right)^2 \text{Var}(\mu_T|\phi) \\ &= \text{Var}(r_{T,T+k}|\mu_T, \phi) + \left(\frac{1 - \beta^k}{1 - \beta}\right)^2 \sigma_u^2 \left(\frac{R^2}{1 - R^2}\right), \end{aligned} \quad (\text{A10})$$

using equation (A7). It follows from equation (6) that

$$\text{Var}(r_{T,T+1}|\mu_T, \phi) = \sigma_u^2. \quad (\text{A11})$$

Combining equations (A10) and (A11) for $k = 1$ gives

$$\text{Var}(r_{T,T+1}|\phi) = \text{Var}(r_{T,T+1}|\mu_T, \phi) + \frac{\sigma_u^2 R^2}{1 - R^2} = \frac{\sigma_u^2}{1 - R^2} = \frac{\text{Var}(r_{T,T+1}|\mu_T, \phi)}{1 - R^2}. \quad (\text{A12})$$

Denote the conditional variance ratio $V_c(k)$ and the unconditional variance ratio $V_u(k)$ as follows:

$$V_c(k) = \frac{(1/k)\text{Var}(r_{T,T+k}|\mu_T, \phi)}{\text{Var}(r_{T+1}|\mu_T, \phi)}; \quad V_u(k) = \frac{(1/k)\text{Var}(r_{T,T+k}|\phi)}{\text{Var}(r_{T,T+1}|\phi)}. \quad (\text{A13})$$

These ratios can then be related as follows, combining (A10), (A12), and (A13):

$$\begin{aligned}
V_u(k) &= \frac{(1/k)\text{Var}(r_{T,T+k}|\phi)(1-R^2)}{\text{Var}(r_{T,T+1}|\mu_T, \phi)} \\
&= \frac{(1/k)\text{Var}(r_{T,T+k}|\mu_T, \phi)(1-R^2)}{\text{Var}(r_{T,T+1}|\mu_T, \phi)} + \frac{1}{k} \left(\frac{1-\beta^k}{1-\beta} \right)^2 R^2 \\
&= (1-R^2)V_c(k) + \frac{1}{k} \left(\frac{1-\beta^k}{1-\beta} \right)^2 R^2.
\end{aligned} \tag{A14}$$

A.5. Permanent and temporary price components in our setting

Fama and French (1988), Summers (1986), and others employ a model in which the log stock price p_t is the sum of a random walk s_t and a stationary component y_t that follows an AR(1) process:

$$p_t = s_t + y_t \tag{A15}$$

$$s_t = \mu + s_{t-1} + \epsilon_t \tag{A16}$$

$$y_t = by_{t-1} + e_t, \tag{A17}$$

where e_t and ϵ_t are mean-zero variables independent of each other, and $|b| < 1$. Noting that $r_{t+1} = p_{t+1} - p_t$, it is easy to verify that equations (A15) through (A17) deliver a special case of our model in equations (1) and (5), in which

$$E_r = \mu \tag{A18}$$

$$\beta = b \tag{A19}$$

$$\mu_t = \mu - (1-b)y_t \tag{A20}$$

$$u_{t+1} = \epsilon_{t+1} + e_{t+1} \tag{A21}$$

$$w_{t+1} = -(1-b)e_{t+1}. \tag{A22}$$

This special case has the property

$$\sigma_{uw} = \text{Cov}(u_{t+1}, w_{t+1}) = -(1-b)\sigma_e^2 < 0, \tag{A23}$$

implying the presence of mean reversion. We also see

$$\sigma_\mu^2 = \text{Var}(\mu_t) = (1-b)^2\sigma_y^2 = (1-b)^2 \frac{\sigma_e^2}{1-b^2} = \frac{1-b}{1+b}\sigma_e^2 \tag{A24}$$

and, therefore, using (21),

$$\text{Cov}(r_{t+1}, r_t) = \beta\sigma_\mu^2 + \sigma_{uw} = \frac{b(1-b)}{1+b}\sigma_e^2 - (1-b)\sigma_e^2 = -\frac{1-b}{1+b}\sigma_e^2 < 0. \tag{A25}$$

Thus, under (A15) through (A17) with $b > 0$, all autocovariances in (21) are negative and all unconditional variance ratios are less than 1.

Table 1
Variance Ratios and Components of Long-Horizon Variance

The first row of each panel reports the ratio $(1/k)\text{Var}(r_{T,T+k}|D_T)/\text{Var}(r_{T+1}|D_T)$, where $\text{Var}(r_{T,T+k}|D_T)$ is the predictive variance of the k -year return based on 206 years of annual data for real equity returns and the three predictors over the 1802–2007 period. The second row reports $\text{Var}(r_{T,T+k}|D_T)$, multiplied by 100. The remaining rows report the five components of $\text{Var}(r_{T,T+k}|D_T)$, also multiplied by 100 (they add up to total variance). Panel A contains results for $k = 25$ years, and Panel B contains results for $k = 50$ years. Results are reported under each of three priors for ρ_{uw} , R^2 , and β . As the prior for one of the parameters departs from the benchmark, the priors on the other two parameters are held at the benchmark priors. The “tight” priors, as compared to the benchmarks, are more concentrated towards -1 for ρ_{uw} , 0 for R^2 , and 1 for β ; the “loose” priors are less concentrated in those directions.

| Prior | ρ_{uw} | | | R^2 | | | β | | |
|--|-------------|-------|-------|-------|-------|-------|---------|-------|-------|
| | Tight | Bench | Loose | Tight | Bench | Loose | Tight | Bench | Loose |
| Panel A. Investment Horizon $k = 25$ years | | | | | | | | | |
| Variance Ratio | 1.30 | 1.36 | 1.26 | 1.31 | 1.36 | 1.15 | 1.42 | 1.36 | 1.34 |
| Predictive Variance | 3.82 | 3.99 | 3.68 | 3.92 | 3.99 | 3.28 | 4.17 | 3.99 | 3.93 |
| IID Component | 2.59 | 2.60 | 2.59 | 2.75 | 2.60 | 2.43 | 2.58 | 2.60 | 2.60 |
| Mean Reversion | -4.13 | -4.01 | -4.10 | -3.04 | -4.01 | -4.51 | -4.28 | -4.01 | -3.97 |
| Uncertain Future μ | 2.91 | 2.86 | 2.84 | 1.70 | 2.86 | 3.51 | 3.14 | 2.86 | 2.79 |
| Uncertain Current μ | 0.97 | 0.96 | 0.94 | 0.75 | 0.96 | 0.92 | 1.17 | 0.96 | 0.93 |
| Estimation Risk | 1.48 | 1.58 | 1.41 | 1.75 | 1.58 | 0.93 | 1.56 | 1.58 | 1.57 |
| Panel B. Investment Horizon $k = 50$ years | | | | | | | | | |
| Variance Ratio | 1.76 | 1.82 | 1.64 | 1.72 | 1.82 | 1.45 | 1.96 | 1.82 | 1.79 |
| Predictive Variance | 5.14 | 5.34 | 4.79 | 5.14 | 5.34 | 4.13 | 5.75 | 5.34 | 5.27 |
| IID Component | 2.59 | 2.60 | 2.59 | 2.75 | 2.60 | 2.43 | 2.58 | 2.60 | 2.60 |
| Mean Reversion | -5.52 | -5.36 | -5.42 | -4.32 | -5.36 | -5.61 | -5.80 | -5.36 | -5.28 |
| Uncertain Future μ | 5.40 | 5.31 | 5.13 | 3.60 | 5.31 | 5.54 | 5.97 | 5.31 | 5.16 |
| Uncertain Current μ | 0.95 | 0.94 | 0.91 | 0.90 | 0.94 | 0.73 | 1.16 | 0.94 | 0.92 |
| Estimation Risk | 1.72 | 1.85 | 1.59 | 2.21 | 1.85 | 1.03 | 1.85 | 1.85 | 1.87 |

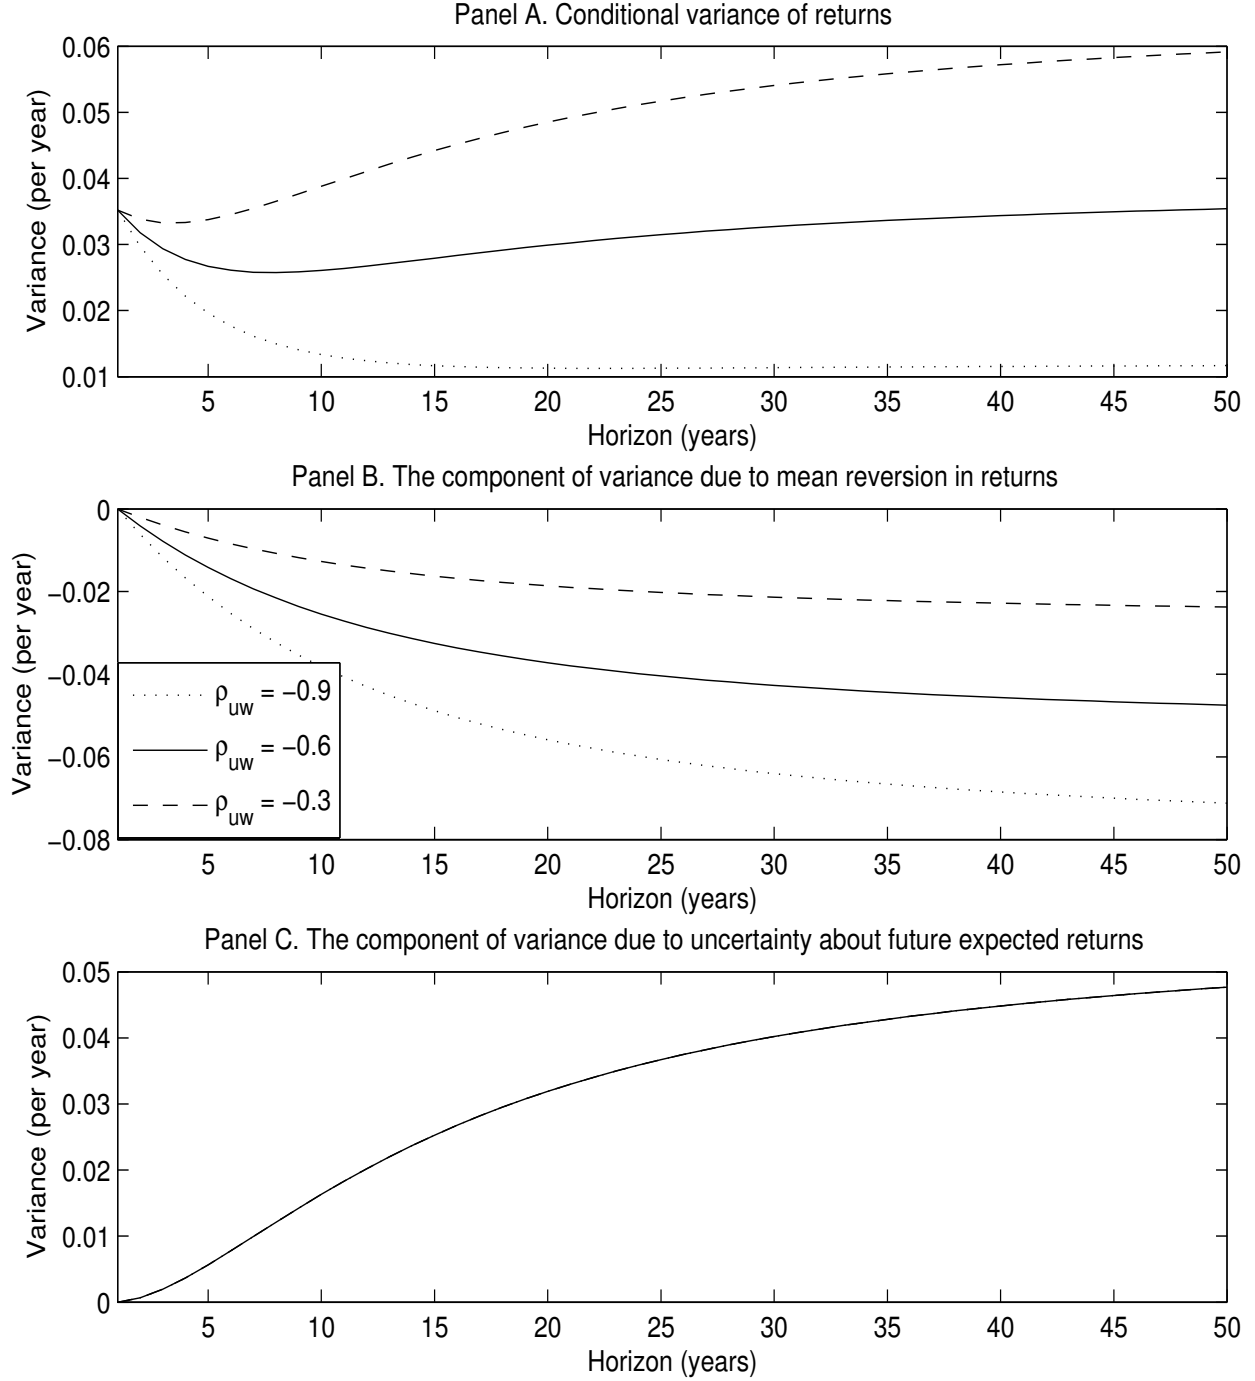


Figure 1. Conditional multiperiod variance and its components for different values of ρ_{uw} . Panel A plots the conditional per-period variance of multiperiod returns from equation (6), $\text{Var}(r_{T,T+k}|\mu_T, \phi)/k$, as a function of the investment horizon k , for three different values of ρ_{uw} . Panel B plots the component of the variance that is due to mean reversion in returns, $\sigma_u^2 2\bar{d}\rho_{uw}A(k)$. Panel C plots the component of this variance that is due to uncertainty about future values of the expected return, $\sigma_u^2 \bar{d}^2 B(k)$. For all three values of ρ_{uw} , variances are computed with $\beta = 0.85$, $R^2 = 0.12$, and an unconditional standard deviation of returns of 20% per year.

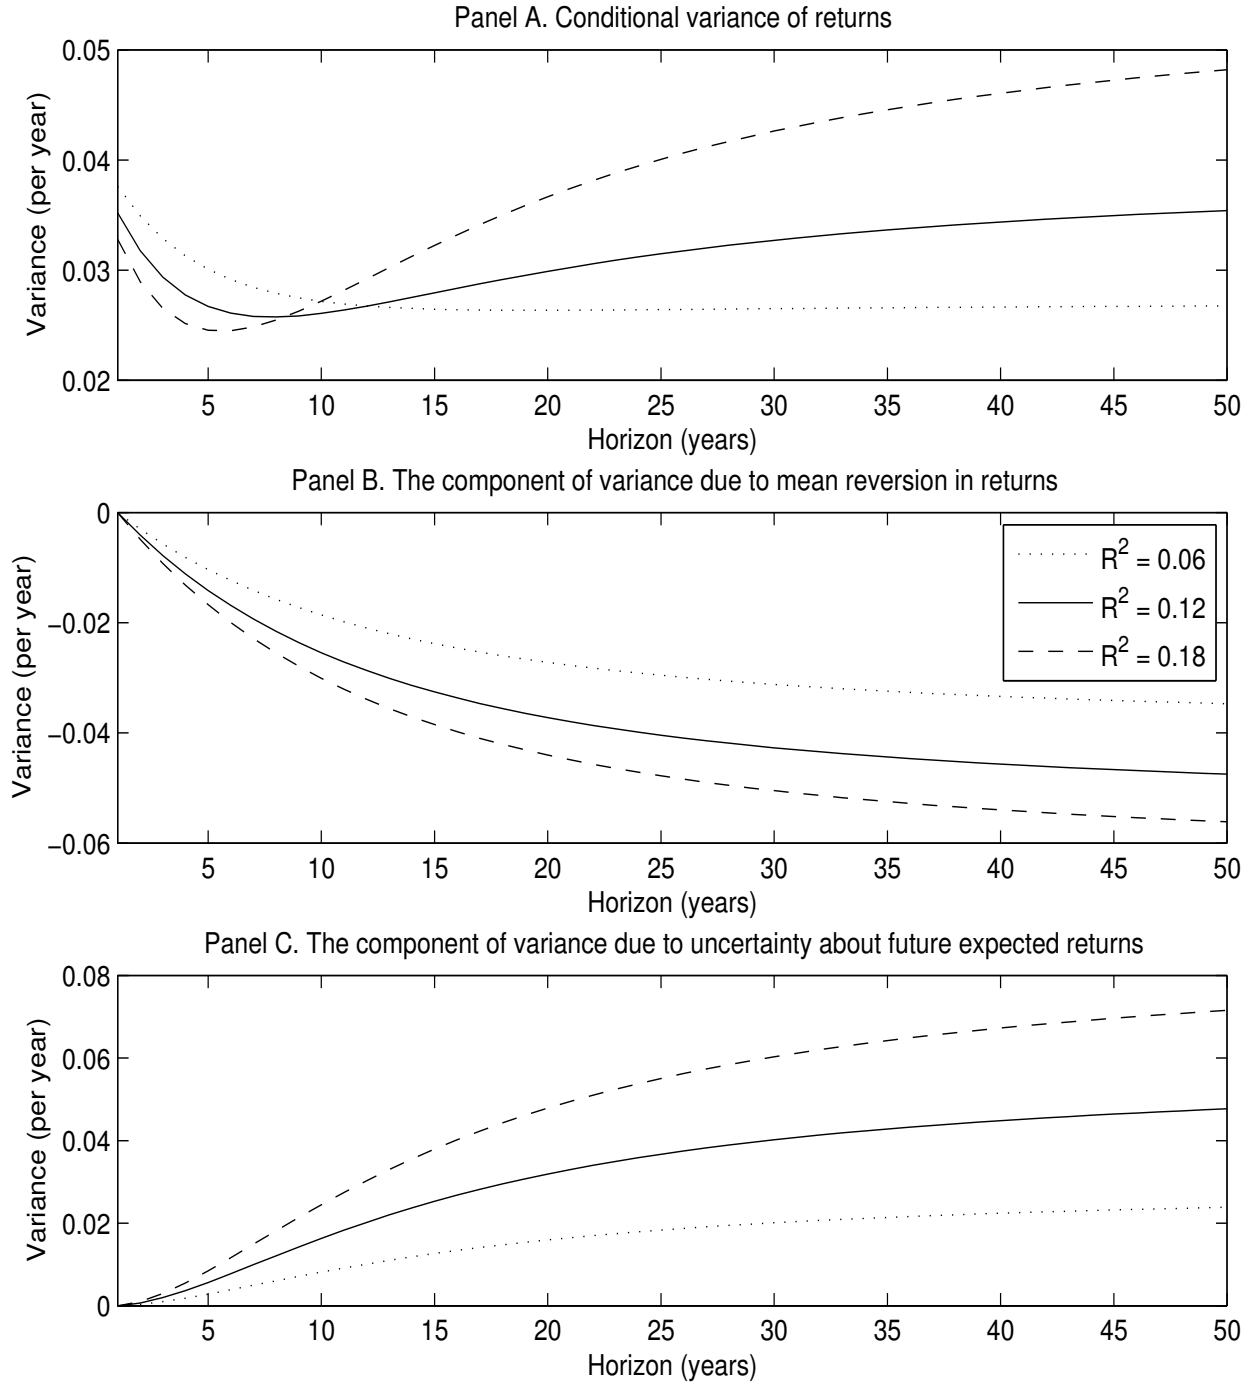


Figure 2. Conditional multiperiod variance and its components for different values of R^2 . Panel A plots the conditional per-period variance of multiperiod returns from equation (6), $\text{Var}(r_{T,T+k}|\mu_T, \phi)/k$, as a function of the investment horizon k , for three different values of R^2 . Panel B plots the component of the variance that is due to mean reversion in returns, $\sigma_u^2 2\bar{d}\rho_{uw}A(k)$. Panel C plots the component of this variance that is due to uncertainty about future values of the expected return, $\sigma_u^2 \bar{d}^2 B(k)$. For all three values of R^2 , variances are computed with $\beta = 0.85$, $\rho_{uw} = -0.6$, and an unconditional standard deviation of returns of 20% per year.




Figure 3. Prior distributions of parameters. The plots display the prior distributions for β , ρ_{uw} , the true R^2 (fraction of variance in the return r_{t+1} explained by the conditional mean μ_t), and the “observed” R^2 (fraction of variance in r_{t+1} explained by the observed predictors x_t). The priors shown for the observed R^2 correspond to the three priors for the true R^2 and the benchmark priors for β and ρ_{uw} .




Figure 4. Posterior distributions of parameters. Panel A plots the posteriors of ρ_{uw} , the correlation between expected and unexpected returns. Panel B plots the posteriors of β , the persistence of the true conditional expected return μ_t . Panel C plots the posteriors of the true R^2 (fraction of variance in the return r_{t+1} explained by μ_t). Panel D plots the posteriors of the “observed” R^2 (fraction of variance in r_{t+1} explained by the observed predictors x_t). The results are obtained by estimating the predictive system on annual real U.S. stock market returns in 1802-2007. Three predictors are used: the dividend yield, the bond yield, and the term spread.

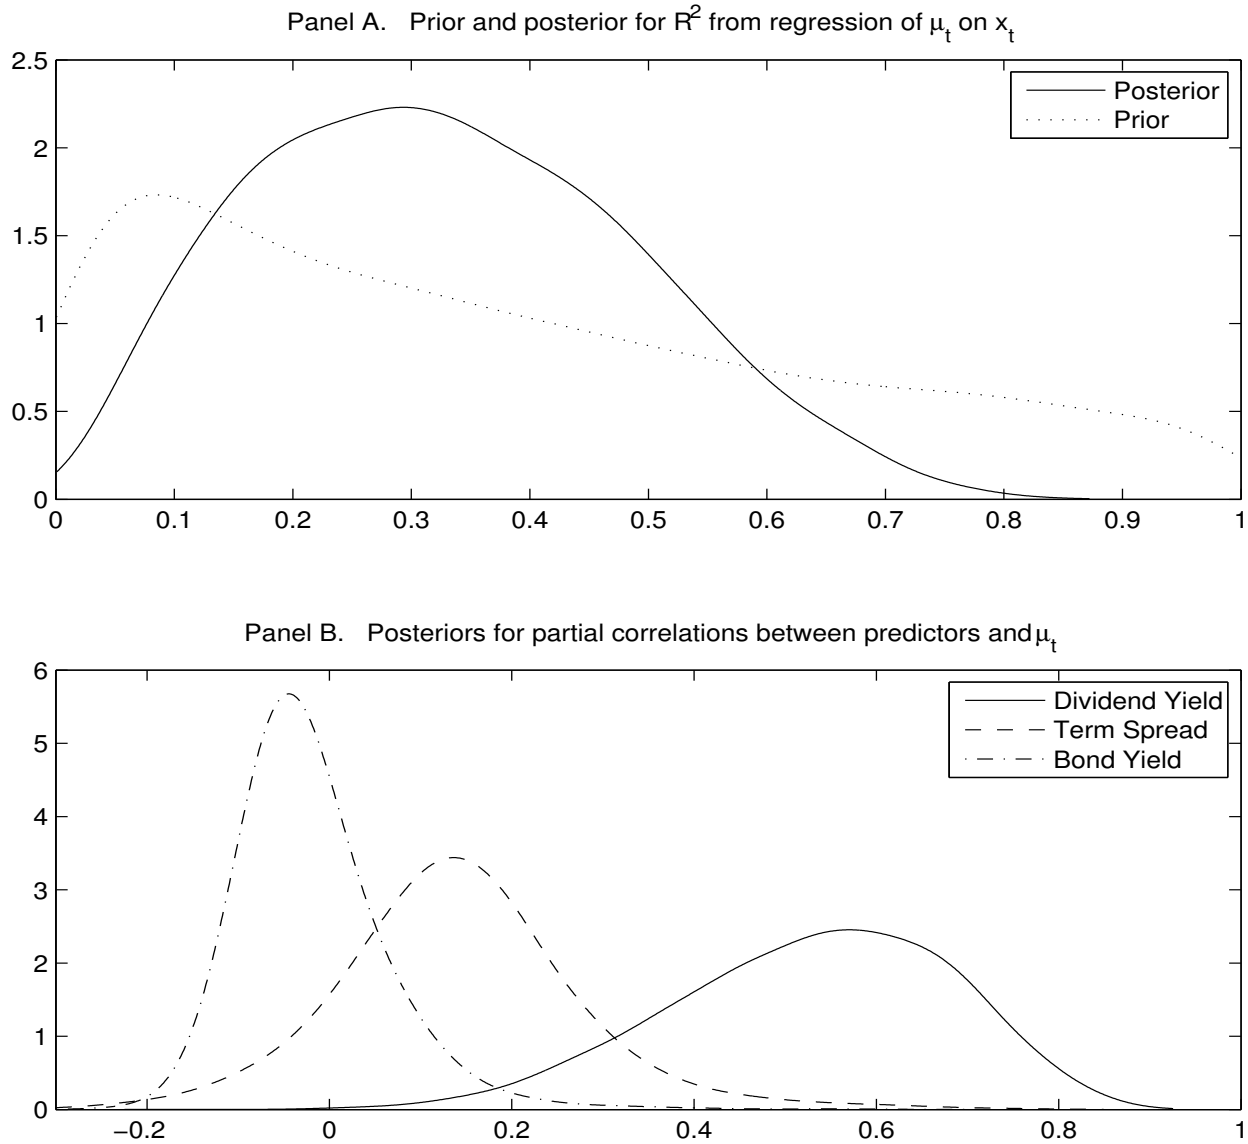


Figure 5. Distributions of parameters related to predictor imperfection. Panel A plots the (implied) prior and posterior of the fraction of variance in the conditional expected return μ_t that can be explained by the predictors. The values smaller than one indicate predictor imperfection. Panel B plots the posteriors of partial correlations between each of the three predictors and μ_t . Benchmark priors are used throughout. The results are obtained by estimating the predictive system on annual real U.S. stock market returns in 1802-2007. Three predictors are used: the dividend yield, the bond yield, and the term spread.

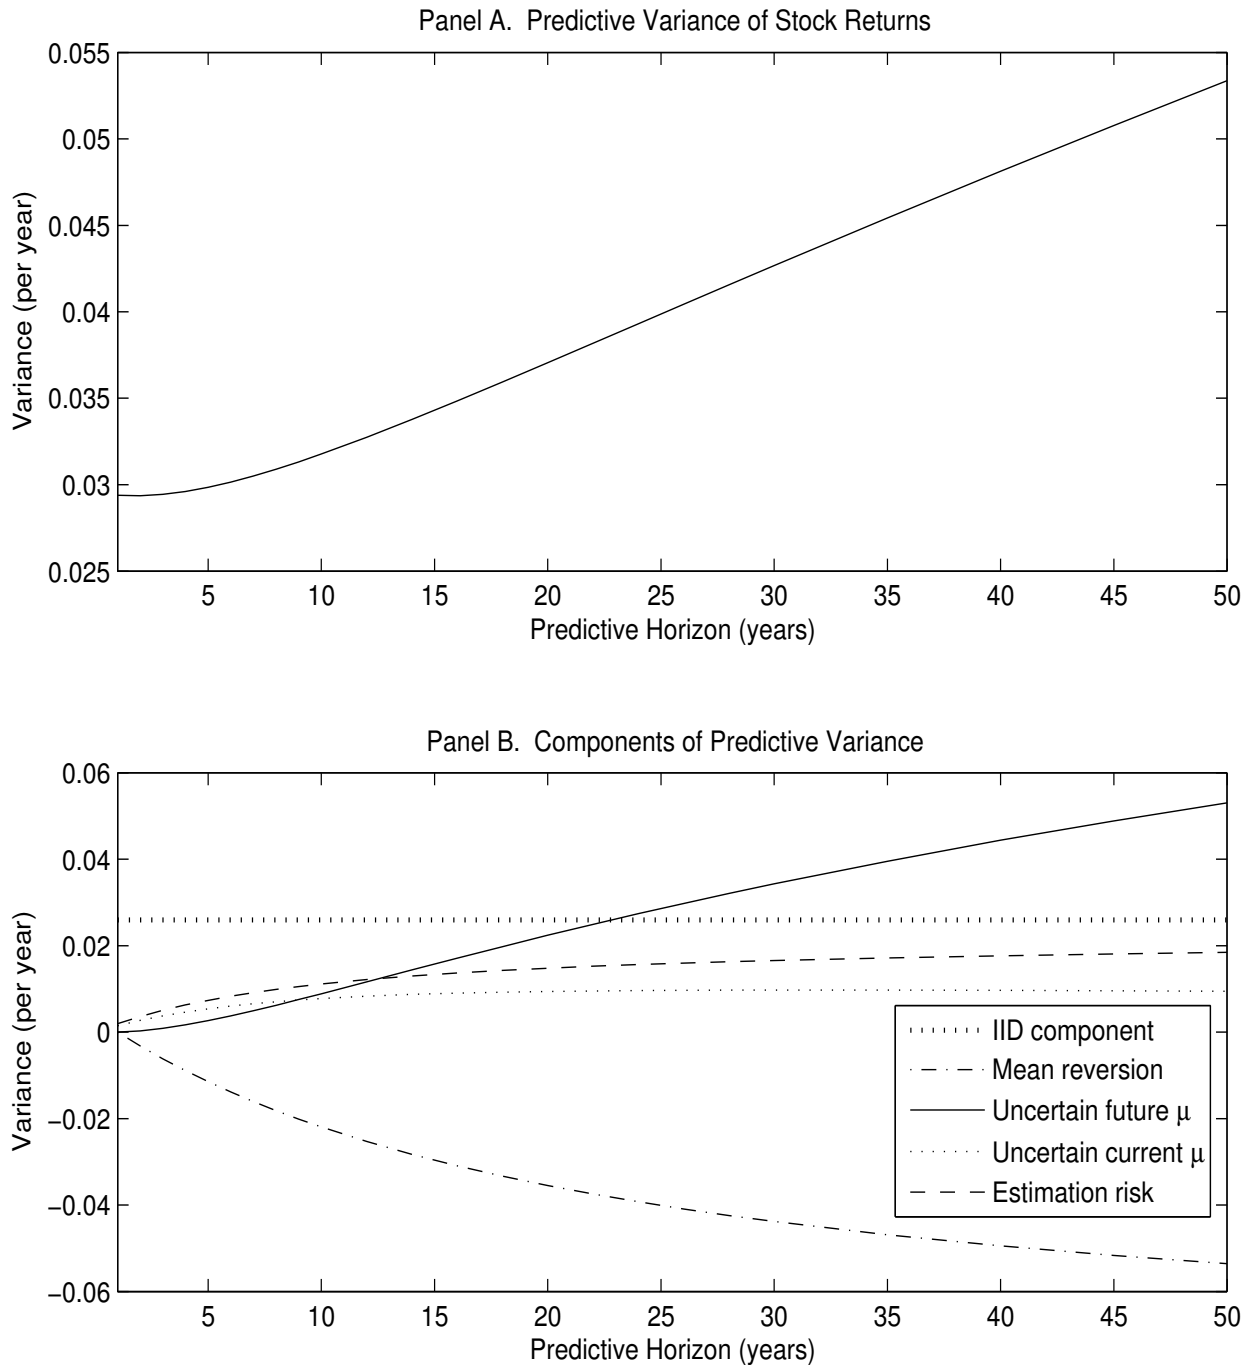


Figure 6. Predictive variance of multiperiod return and its components. Panel A plots the variance of the predictive distribution of long-horizon returns, $\text{Var}(r_{T,T+k}|D_T)$. Panel B plots the five components of the predictive variance. All quantities are divided by k , the number of periods in the return horizon. The results are obtained by estimating the predictive system on annual real U.S. stock market returns in 1802-2007. Three predictors are used: the dividend yield, the bond yield, and the term spread.




Figure 7. Priors and posteriors for predictor imperfection. The plots display prior and posterior distributions under the predictive system (System 2) in which expected return depends on a vector of observable predictors, x_t , as well as a missing predictor, π_t , that obeys an AR(1) process. The top panels display prior distributions for σ_π , the standard deviation of π_t , under different degrees of predictor imperfection. The middle panels display the corresponding posteriors for σ_π . The bottom panels display the posterior distributions of ΔR^2 , the “true” R^2 for one-period returns when conditioning on $\{x_t, \pi_t\}$ minus the “observed” R^2 when conditioning only on x_t . The left-hand panels are based on annual data from 1802–2007 for real U.S. stock returns and three predictors: the dividend yield, the bond yield, and the term spread. The right-hand panels are based on quarterly data from 1952Q1–2006Q4 for real returns and three predictors: the dividend yield, CAY, and the bond yield.




Figure 8. Predictive variance under predictor imperfection. The plots display predictive variance under the predictive system (System 2) in which expected return depends on a vector of observable predictors, x_t , as well as a missing predictor, π_t , that obeys an AR(1) process. Predictive variances are shown for the two imperfect-predictor cases as well for the case of perfect predictors ($\sigma_{\pi} = 0$). Panel A is based on annual data from 1802–2007 for real U.S. stock returns and three predictors: the dividend yield, the bond yield, and the term spread. Panel B is based on quarterly data from 1952Q1–2006Q4 for real returns and three predictors: the dividend yield, CAY, and the bond yield.




Figure 9. Predictive variance under model uncertainty. The plot displays predictive variance computed using the model-uncertainty framework of Avramov (2002). The predictive variance includes uncertainty about which, if any, of a set of predictors enter a predictive regression for returns. The top panel is based on annual data from 1802–2007 for real U.S. stock returns and three predictors: the dividend yield, the bond yield, and the term spread. The bottom panel is based on quarterly data from 1952Q1–2006Q4 for real returns and three predictors: the dividend yield, CAY, and the bond yield.




Figure 10. Sample variance ratios of annual real equity returns, 1802–2007. The plot displays the sample variance ratio $\hat{V}(k) = \hat{\text{Var}}(r_{t,t+k}) / (k\hat{\text{Var}}(r_{t,t+1}))$, where $\hat{\text{Var}}(r_{t,t+k})$ is the unbiased sample variance of k -year log returns, computed at an overlapping annual frequency. Also shown are the 1st, 10th, and 50th percentiles of the Monte Carlo sampling distribution of $\hat{V}(k)$ under the hypothesis that annual log returns are independently and identically distributed as normal.

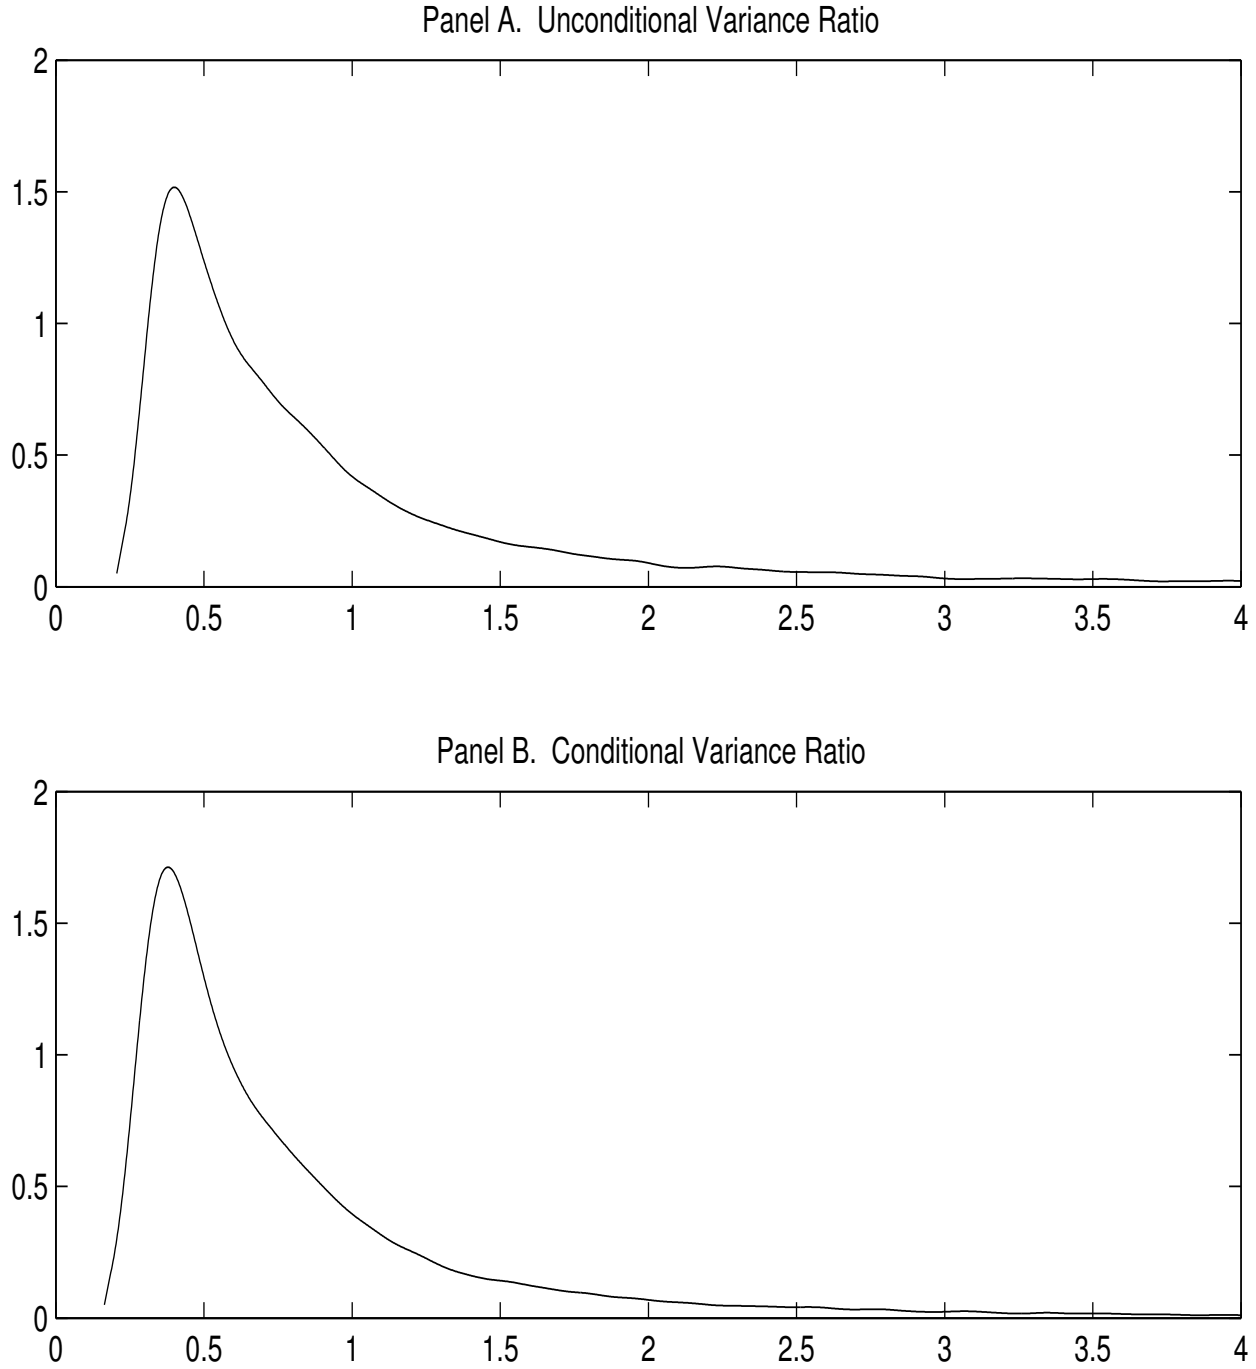


Figure 11. Posterior distributions for 30-year variance ratios. Panel A plots the posterior distribution of the unconditional variance of 30-year stock market returns, $\text{Var}(r_{T,T+30}|\phi)$, divided by 30 times the unconditional variance of one-year returns, $\text{Var}(r_{T+1}|\phi)$. Panel B plots the analogous ratio for the conditional variance, $\text{Var}(r_{T,T+30}|D_T, \phi)$. (The posterior mean of that variance is the first term of the predictive variance in equation (25).) The results are obtained by estimating System 1 on annual real U.S. stock market returns in 1802-2007. Three predictors are used: the dividend yield, the bond yield, and the term spread.

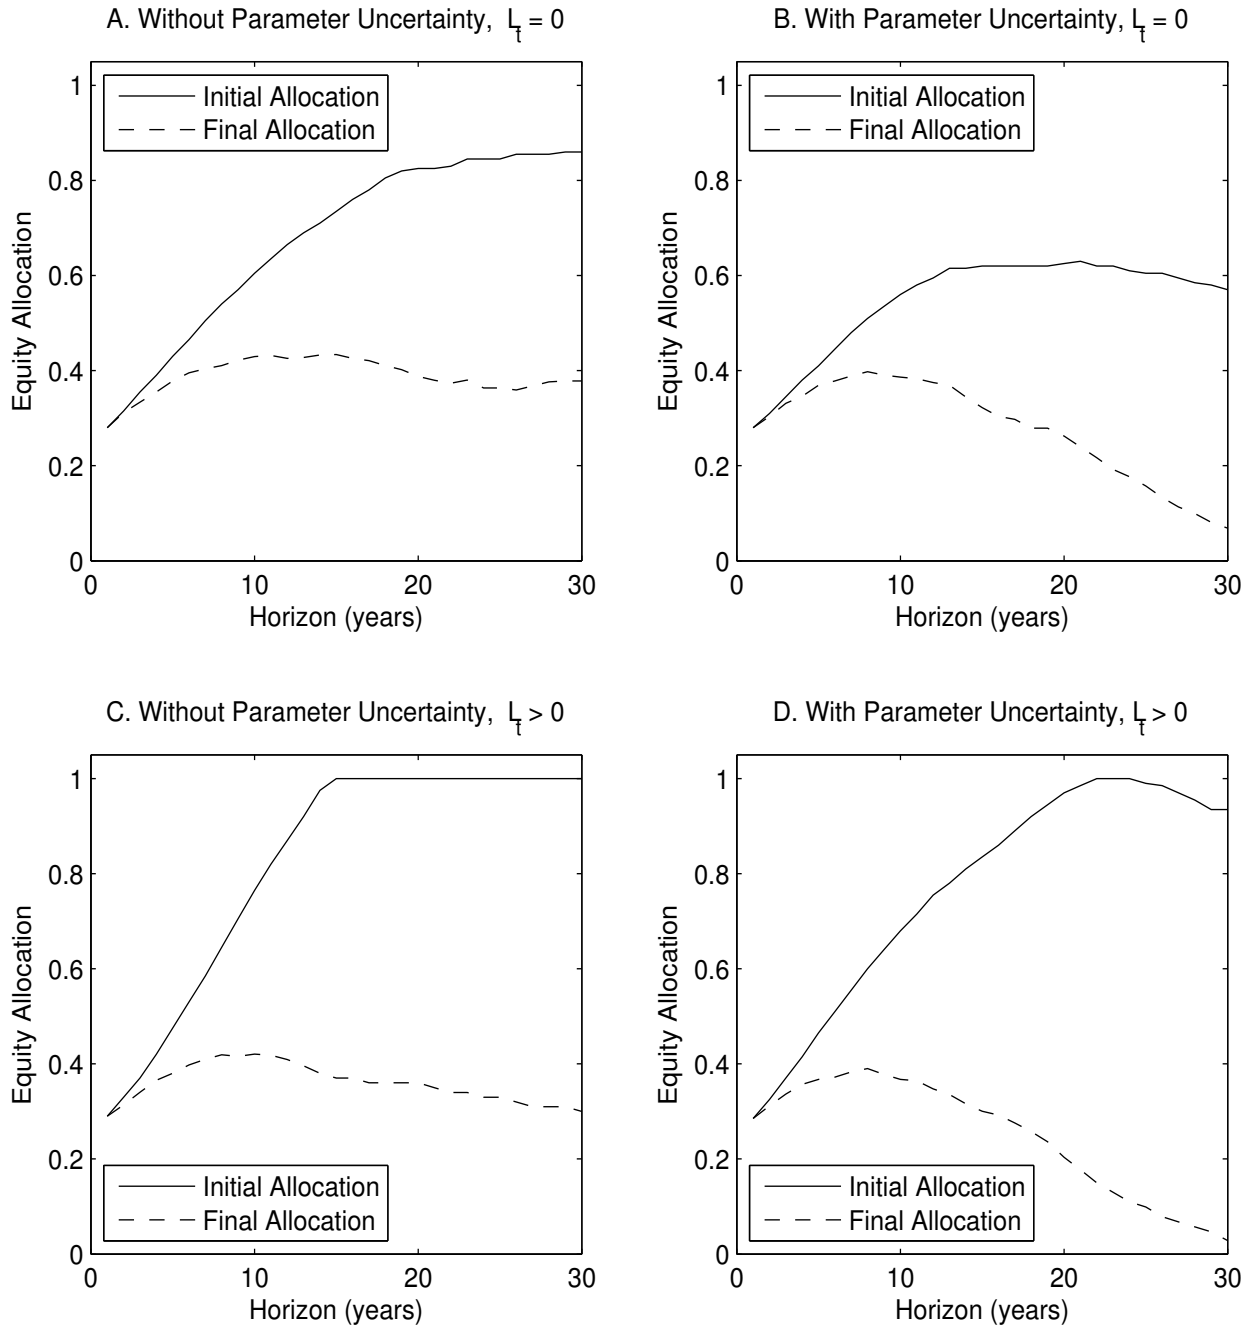


Figure 12. Parameter uncertainty and target-date funds. The figure plots equity allocations w_1 (solid line) and w_K (dashed line) for a long-horizon investor with utility for end-of-horizon wealth (W) given by $W^{1-A}/(1-A)$. At the beginning of a K -period horizon, the investor commits to a strategy in which the equity allocation evolves linearly from the first-period allocation w_1 to the final-period allocation w_K . The remaining portion of the investor's portfolio is allocated to a riskless asset, assumed to provide a constant real return of 2% per year. Relative risk aversion (A) equals 8. The investor chooses both w_1 and w_K on the interval $(0, 1)$ to maximize expected utility. The investor incorporates parameter uncertainty in Panels B and D but not in Panels A and C. The investor has no labor income savings in Panels A and B (equation (30)). In Panels C and D, he does save from labor income, and his wealth evolves as in equation (31).

References

- Ang, Andrew, and Monika Piazzesi, 2003, A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables, *Journal of Monetary Economics* 50, 745–787.
- Avramov, Doron, 2002, Stock return predictability and model uncertainty, *Journal of Financial Economics* 64, 423–458.
- Balduzzi, Pieluigi, and Anthony W. Lynch, 1999, Transaction costs and predictability: some utility cost calculations, *Journal of Financial Economics* 52, 47–78.
- Barberis, Nicholas, 2000, Investing for the long run when returns are predictable, *Journal of Finance* 55, 225–264.
- Ben-David, Itzhak, John R. Graham, Campbell R. Harvey, 2010, Managerial Miscalibration, Working paper, Ohio State University and Duke University.
- Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Goldstein, 2007, Portfolio choice over the life-cycle when the stock and labor markets are cointegrated, *Journal of Finance* 62, 2123–2167.
- Bodie, Zvi, Robert C. Merton, and William F. Samuelson, 1992, “Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model,” *Journal of Economic Dynamics and Control* 16, 427–449.
- Brandt, Michael W., Amit Goyal, Pedro Santa-Clara, and Jonathan Stroud, 2005, A simulation approach to dynamic portfolio choice with an application to learning about return predictability, *Review of Financial Studies* 18, 831–873.
- Brandt, Michael W., and Qiang Kang, 2004, On the relationship between the conditional mean and volatility of stock returns: A latent VAR approach, *Journal of Financial Economics* 72, 217–257.
- Brandt, Michael W., Pedro Santa-Clara, and Rossen Valkanov, 2009, Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns, *Review of Financial Studies* 22, 3411–3447.
- Campbell, John Y., 1991, A variance decomposition for stock returns, *The Economic Journal* 101, 157–179.
- Campbell, John Y., Yeung Lewis Chan and Luis M. Viceira, 2003, A multivariate model of strategic asset allocation, *Journal of Financial Economics* 67, 41–80.
- Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997, *The Econometrics of Financial Markets* (Princeton University Press, Princeton, NJ).
- Campbell, John Y., and Luis M. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors* (Oxford University Press).
- Campbell, John Y., and Luis M. Viceira, 2005, The term structure of the risk-return tradeoff, *Financial Analysts Journal* 61.
- Chacko, George, and Luis M. Viceira, 2005, Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets, *Review of Financial Studies* 18, 1369–1402.
- Cocco, Joao F., Francisco J. Gomes, and Pascal J. Maenhout, 2005, “Consumption and Portfolio Choice over the Life Cycle,” *Review of Financial Studies* 18, 491–533.

- Conrad, Jennifer, and Gautam Kaul, 1988, Time-variation in expected returns, *Journal of Business* 61, 409-425.
- Cremers, K.J. Martijn, Stock return predictability: A Bayesian model selection perspective, *Review of Financial Studies* 15, 1223–1249.
- Dangl, Thomas, and Michael Halling, 2006, Equity return prediction: Are coefficients time-varying?, Working paper, University of Vienna.
- Detemple, Jérôme, René Garcia, and Marcel Rindisbacher, 2005, Intertemporal asset allocation: A comparison of methods, *Journal of Banking and Finance* 29, 2821–2848.
- Duffee, Gregory R., 2006, Are variations in term premia related to the macroeconomy? Working paper, University of California, Berkeley.
- Fama, Eugene F., and Kenneth R. French, 1988, Permanent and temporary components of stock prices, *Journal of Political Economy* 96, 246–273.
- Fama, Eugene F., and Kenneth R. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23–49.
- Gomes, Francisco J., Laurence J. Kotlikoff, and Luis M. Viceira, 2008, “Optimal Life-Cycle Investing with Flexible Labor Supply: A Welfare Analysis of Life-Cycle Funds,” *American Economic Review: Papers and Proceedings* 98, 297–303.
- Gomes, Francisco, and Alexander Michaelides, 2005, “Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence,” *Journal of Finance* 60, 869–904.
- Gordon, Katherine D. and Kimberly A. Stockton, 2006, *Funds for Retirement: The ‘Life-Cycle’ Approach* (The Vanguard Group, Valley Forge, PA).
- Greer, Boyce, 2004, *The case for age-based life-cycle investing* (Fidelity Investments, Boston, MA).
- Hansen, Lars P., 2007, Beliefs, doubts and learning: Valuing macroeconomic risk, *American Economic Review P&P* 97, 1–30.
- Heaton, John, and Deborah Lucas, 2000, Portfolio choice in the presence of background risk, *The Economic Journal* 110, 1–26.
- Hoevenaars, Roy, Roderick Molenaar, Peter C. Schotman, and Tom Steenkamp, 2007, Strategic asset allocation for long-term investors: Parameter uncertainty and prior information, Working paper, Maastricht University.
- Investment Company Institute, 2011, *The U.S. Retirement Market: Third Quarter 2010*. Washington, DC.
- Johannes, Michael, Nicholas Polson, and Jon Stroud, 2002, Sequential optimal portfolio performance: Market and volatility timing, Working paper, Columbia University.
- Kim, M., Nelson, C. and Startz, R., 1991, Mean reversion in stock prices? A reappraisal of the empirical evidence, *Review of Economic Studies* 58, 515-28.
- Lamoureux, Christopher G., and Guofu Zhou, 1996, Temporary components of stock returns: What do the data tell us? *Review of Financial Studies* 9, 1033–1059.

- Lettau, Martin, and Sydney C. Ludvigson, 2001, Consumption, aggregate wealth, and expected stock returns, *Journal of Finance* 56, 815–849.
- Lo, Andrew and MacKinlay, Craig A., 1988, Stock market prices do not follow random walks: Evidence from a simple specification test, *Review of Financial Studies* 1, 41–66.
- Lo, Andrew and MacKinlay, Craig A., 1989, The size and power of the variance ratio test in finite samples, *Journal of Econometrics* 40, 203–238.
- Lynch, Anthony W., 2001, Portfolio choice and equity characteristics: characterizing the hedging demands induced by return predictability, *Journal of Financial Economics* 62, 671–730.
- Lynch, Anthony W., and Pierluigi Balduzzi, 2000. Transaction costs and predictability: the impact on portfolio choice, *Journal of Finance* 55, 2285–2310.
- Lynch, Anthony W., and Sinan Tan, 2009, Labor Income Dynamics at Business-cycle Frequencies: Implications for Portfolio Choice, *Journal of Financial Economics*, forthcoming.
- Pástor, Ľuboš, and Robert F. Stambaugh, 2009, Predictive systems: Living with imperfect predictors, *Journal of Finance* 64, 1583–1628.
- Poterba, James M. and Lawrence H. Summers, 1988, Mean reversion in stock returns: Evidence and implications, *Journal of Financial Economics* 22, 27–60.
- Richardson, Matthew, 1993, Temporary components of stock prices: a skeptic's view, *Journal of Business and Economic Statistics* 11, 199–207.
- Richardson, Matthew, and James H. Stock, 1989, Drawing inferences from statistics based on multiyear asset returns, *Journal of Financial Economics* 25, 323–348.
- Rytchkov, Oleg, 2007, Filtering out expected dividends and expected returns, Working paper, MIT.
- Schotman, Peter C., Rolf Tschernig, and Jan Budek, 2008, Long memory and the term structure of risk, *Journal of Financial Econometrics*, forthcoming.
- Siegel, Jeremy J., 1992, The real rate of interest from 1800–1990: A study of the U.S. and the U.K., *Journal of Monetary Economics* 29, 227–252.
- Siegel, Jeremy J., 2008, *Stocks for the Long Run*, 4th ed. (McGraw Hill, New York, NY).
- Summers, Lawrence H., 1996, Does the stock market rationally reflect fundamental values? *Journal of Finance* 41, 591–601.
- Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375–421.
- Viceira, Luis M., 2001, “Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income,” *Journal of Finance* 56, 433–470.
- Viceira, Luis M., 2008, Life-cycle funds, forthcoming in Annamaria Lusardi, ed., *Overcoming the Saving Slump: How to Increase the Effectiveness of Financial Education and Saving Programs* (University of Chicago Press, Chicago, IL).
- Wachter, Jessica, 2010, Asset allocation, *Annual Review of Financial Economics*, forthcoming.