Long-Term Natural Course of Small Nonfunctional Pancreatic Neuroendocrine Tumors in MEN1-Results From the Dutch MEN1 Study Group

Pieterman, Carolina R.C.; de Laat, Joanne M.; Twisk, Jos W.R.; van Leeuwaarde, Rachel S.; de Herder, Wouter W.; Dreijerink, Koen M.A.; Hermus, Ad R.M.M.; Dekkers, Olaf M.; van der Horst-Schrivers, Anouk N.A.; Drent, Madeleine L.; Bisschop, Peter H.; Havekes, Bastiaan; Borel Rinkes, Inne H.M.; Vriens, Menno R.; Valk, Gerlof D.

published in
Journal of Clinical Endocrinology and Metabolism
2017

DOI (link to publisher)
10.1210/jc.2017-00372

document version
Publisher's PDF, also known as Version of record

document license
Article 25fa Dutch Copyright Act

JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 52, No. 5, Oct. 2017, pp. 2277–2303

doi:10.1017/S002210901700059X

Long-Term versus Short-Term Contingencies in Asset Allocation

Mahmoud Botshekan and André Lucas*

Abstract

We investigate whether long-term and short-term components of typical conditioning variables in asset pricing studies, such as the dividend yield or yield spread, have different implications for optimal asset allocation. We argue that short-term components relate mostly to momentum, and long-term components relate mostly to mean-reversion effects, respectively. Therefore, they may have a different information content for investors with different horizons. We obtain improvements in terms of out-of-sample Sharpe ratios and expected utilities for decomposed state variables that directly reflect information related to the stock market, such as the dividend yield and stock market trend.

I. Introduction

There is now a large body of literature claiming that asset returns are predictable to some extent.[1] The existence of a partly predictable component in asset returns has profound implications for optimal asset-allocation decisions. In particular, if we can capture changes in investment opportunity sets by a (small) set of state variables, optimal asset allocations also become functions of these state variables. Empirical studies have tried to exploit this feature by using different

*Botshekan (corresponding author), m.botshekan@ase.ui.ac.ir, Faculty of Administrative Sciences and Economics, University of Isfahan and University College Dublin; Lucas, a.lucas@vu.nl, School of Business and Economics, Vrije Universiteit Amsterdam, and Tinbergen Institute. We thank an anonymous referee, Stephen Brown (the editor), Frank de Jong, Ralph Koijen, Philip Stork, and Pieter Jelle van der Sluis for detailed and helpful comments on earlier versions of this article, as well as participants at the 2012 European Finance Association meeting and seminar participants at the Duisenberg School of Finance. Lucas thanks the Dutch Science Foundation (NWO) under Grant VICI-453-09-005 for financial support. Botshekan acknowledges the support of Science Foundation Ireland under Grant 08/SRC/FMC1389 and the support from the Iran Ministry of Science, Research, and Technology. For programming, we used the programming package OxMetrics; see Doornik (2009) and Doornik and Ooms (2007) for further information.

[1]See, for example, Cochrane (1999), Campbell and Viceira (1999), Lettau and Ludvigson (2005), Ang and Bekaert (2007), Cochrane (2008), and Van Binsbergen and Koijen (2010), among many others.

approaches to model the dependence of asset-allocation decisions on candidate state variables, such as dividend–price ratios and yield spreads.[2]

In this article, we go one step further and investigate whether investors can benefit from reacting differently to short-term versus long-term information in commonly used state variables. This is particularly important in the context of tactical and strategic asset-allocation decisions: Different components of state variables may capture different dynamics in investment opportunity sets. For instance, many studies demonstrate that momentum and mean reversion in stock returns have important yet different implications for the asset-allocation decisions of short-term versus long-term investors.[3] In a similar vein, we hypothesize that long-term and short-term variation in state variables may correlate with different types of changes in investment opportunity sets and may, therefore, have a different information content for different investors. Short-term components may capture the short-term predictability of returns in excess of expected returns, for example, relating to such phenomena as momentum or overreaction. Long-term components, conversely, may capture secular trends in expected returns and long-term predictability related to, for example, mean reversion and persistent changes in fundamentals.[4] As a result and depending on his or her horizon, an investor might react differently to an increase in a particular state variable (e.g., the dividend yield) depending on whether the increase is due to a change in the short-term or the long-term component of that variable.

To investigate how short-term and long-term dynamics in state variables affect portfolio choice, we use standard techniques for state-variable decomposition from the macro literature and apply these to the asset-allocation context. We use a flexible, semiparametric approach from the finance literature to link the state variables to concrete asset allocations. Our approach thus combines two lines of literature: i) the filtering techniques from the macroeconomic and econometric

[2]For this purpose, some studies consider parametric models with fixed parameters such as vector autoregressions (VARs) to model the relation between time-varying investment opportunities and state variables; see, for example, Lynch (2001) and Campbell, Chan, and Viceira (2003). Other studies use various Bayesian approaches to deal with parameter and model uncertainty/instability, for example, Kandel and Stambaugh (1996), Barberis (2000), Cremers (2002), Wachter and Warusawitharana (2009), Pettenuzzo and Timmermann (2011), and Guidolin and Timmermann (2007). An alternative line of approach is the nonparametric approach of Brandt (1999) and Ait-Sahalia and Brandt (2001) adopted in this paper.

[3]See Balvers and Mitchell (1997) and Brandt (1999) for the effect of momentum and Campbell and Viceira (1999), Wachter (2002), and Campbell et al. (2003) for the effect of mean reversion on optimal asset allocation. Koijen, Rodriguez, and Sbuelz (2009) investigate the joint presence of momentum and mean reversion in asset prices and show that momentum strongly affects the short-term allocations to stocks, whereas mean reversion in equity returns is the dominant force for longer investment horizons.

[4]In this sense, our results also contribute to the debate about the source of predictability of financial ratios such as the dividend yield. Time variation in expected returns captured by dividend yields or other variables can be attributed to shifts in underlying risk (e.g., low dividend yields may imply low subsequent discount rates) or to irrational mispricing (e.g., low dividend yields may imply stocks are overpriced). Ferson and Harvey (1991), (1993), (1998), for example, decompose the predictability associated with conditional variables into a part that can be explained by a rational asset pricing model (shifts or secular trends in risk premia or factor loadings) and a part that is related to market inefficiencies. The relative importance of these sources of predictability in asset-allocation decisions is examined by Tamayo (2002) and Shanken and Tamayo (2012) using a Bayesian approach. Our results also contribute to this debate.

literature to decompose time series into their long-term and short-term components, and ii) the semiparametric approach of Brandt (1999) and Ait-Sahalia and Brandt (2001) to determine the relative importance of different state variables for asset-allocation decisions. The multiple-state-variable approach of Ait-Sahalia and Brandt (2001) uses an index of different state variables to obtain a semiparametric estimate of the optimal asset-allocation decisions of an expected utility-maximizing investor. The semiparametric nature of the specification allows for much flexibility. It also enables us to exploit the first-order conditions of the asset-allocation problem inside a generalized method of moments (GMM) framework to determine the importance of the different state-variable components. The long-term and short-term components are used as (new) state variables in the original asset-allocation problem of Ait-Sahalia and Brandt (2001). The novelty in our approach thus lies in decomposing commonly used state variables into their long-term and short-term components, rather than in proposing new state variables per se.

Decomposing time series into long-term and short-term components has a long history in time-series analysis, particularly in macroeconomics. Widely varying techniques have been proposed in the literature, including the familiar filter of Hodrick and Prescott (1997), newer bandpass filters like those of Christiano and Fitzgerald (2003), and approaches based on unobserved components models like those of Harvey (1990) and Harvey and Jaeger (1993). Each of these methods has its own advantages and drawbacks. In the current article, we focus on the more recent filtering techniques of Christiano and Fitzgerald (2003). The main advantage of these techniques is that they allow us to be explicit regarding what we define as the short-term or cyclical component of a state variable. In particular, the notion of short term may have a different connotation for different investors and, as such, may depend on the investment horizon of the decision maker. The Christiano–Fitzgerald filter allows us to explicitly investigate the sensitivity of investment decisions to the definition of the short-term component. We run this filter recursively to avoid any potential look-ahead biases.

As a side contribution of the current article, we extend the approach of Ait-Sahalia and Brandt (2001) to account for short-sale constraints. The inclusion of short-sale constraints is important to prevent unrealistic asset allocations and extreme leverage. However, short-sale restrictions transform the standard asset-allocation problem into a constrained optimization problem. As a result, the first-order conditions for the optimal asset allocation should be augmented with the appropriate Lagrange–Kuhn–Tucker conditions and multipliers. If these are omitted, the standard first-order conditions are incorrect. Because the first-order conditions are the prime ingredient for the moment conditions in the subsequent GMM stage, their omission may lead to a biased or inefficient inference on the relative importance of long-term versus short-term components for asset allocation. We illustrate the potential effects of this and explain how the multipliers can easily be included in the analysis to avoid this problem.

We apply our approach to a portfolio-choice problem involving three assets: stocks, bonds, and a risk-free asset. We use U.S. data for the period Jan. 1954–Dec. 2012. For an investor with a 1-month investment horizon, we find that the relative importance of long-term versus short-term components depends on the state

https://doi.org/10.1017/S002210901700059X Published online by Cambridge University Press

variable under consideration. For such state variables as the dividend yield, stock market trend, and default spread, the short-term components of the state variables receive a relatively larger weight in the asset-allocation decisions than their long-term counterparts. For such state variables as the term spread and dividend growth rate, the long-term components dominate the asset-allocation decisions. For the term spread, the optimal index weights yield a result that is close to that when using the nondecomposed state variables.

To check the economic significance of our results, we implement the induced investment strategies in an out-of-sample backtesting framework. For this purpose, we first estimate the optimal weights of the long-term and short-term components over the period Jan. 1963–Dec. 2002. Using these estimates, we then implement the asset-allocation strategies using the realized value of the long-term and short-term components of the state variables over the subsequent period Jan. 2003–Dec. 2012. Using a holding period of 1 month, the portfolio returns are stored and used to compute average returns, Sharpe ratios, and average utilities. We also check how the resulting strategies perform compared with similar strategies that consider only the nondecomposed state variable.

We find improvements in terms of out-of-sample Sharpe ratios and certainty equivalents for such decomposed state variables as the dividend yield and stock market trend. The improvements can be mainly attributed to higher average returns rather than to lower risk levels. The results for these two state variables can be well connected to the finance literature. The dividend yield is one of the most widely mentioned candidates when discussing stock return predictability (see, e.g., Barberis (2000), Campbell et al. (2003), Cochrane (2008), Lettau and Van Nieuwerburgh (2008), Wachter and Warusawitharana (2009), and Shanken and Tamayo (2012)). The stock market trend can be interpreted as a proxy for the momentum effect, which is also well documented (see, e.g., Jegadeesh and Titman (1993) and Lewellen (2002)). Moreover, given the relative persistence of dividends compared with prices, the short-term component of dividend yields can be regarded as another proxy for momentum. Following Koijen et al. (2009), these short-term components should therefore receive substantial weights for investors with a short (1-month) horizon.

We subject our results to a range of robustness checks. We find that the results are not sensitive to using different decomposition techniques, such as the Hodrick–Prescott filter. The results also persist if we consider different risk-aversion levels in the asset-allocation problem. We also study the effect of longer investment horizons of 3 months up to 1 year. Allocations based on decomposed state variables continue to outperform allocations based on nondecomposed state variables, in particular for the dividend yield and stock market trend. We find that for longer-term investors, the long-term component of state variables such as the dividend yield and default spread plays a more important role in the asset-allocation decision. For such state variables as the term spread and stock market trend, the horizon appears to have less of an effect. This is in line with Koijen et al. (2009), who use the dividend yield to capture the mean reversion in stock prices. Many studies have found that the investor's investment horizon has a significant impact on optimal asset allocations; see, for example, Brennan, Schwartz, and Lagnado (1997) and Barberis (2000). Our article takes a new perspective

https://doi.org/10.1017/S002210901700059X Published online by Cambridge University Press

on this topic by investigating how the decomposition of state variables into their long-term and short-term components can help investors with different horizons to efficiently incorporate information relating to different frequencies.

Economically, our approach also provides a new perspective on the debate about the sources of predictability of financial ratios such as the dividend yield that can be attributed either to shifts in underlying risk premia or to irrational mispricing (see also footnote 4). Given the relative persistence of dividends compared with stock prices, our short-term component in the dividend yield is a good candidate for capturing mispricing effects (overreaction or momentum), whereas the long-term component may capture rational and more secular shifts in risk premia. The relative importance of our short-term versus long-term component may then reflect the relative importance of each of these two sources of return predictability for investors with different horizons. The results show that investors with a 1-month horizon put more weight on the short-term component (mispricing), whereas investors with a 1-year horizon are more concerned with the long-term component (secular shifts in risk premia).

The remainder of the article is organized as follows: In Section II, we introduce the methodology for estimating the optimal portfolio weights under short-sale constraints and explain the filtering techniques used for the decomposition of state variables into their short-term and long-term components. Section III describes the data. Section IV presents the empirical estimation results. In Section V, we test the effectiveness of induced investment strategies in an out-of-sample backtesting framework. In Section VI, we discuss the robustness checks. Section VII concludes. A number of more technical issues and supplementary results are gathered in the Internet Appendix (available at www.jfqa.org).

II. Methodology

In this section, we recapitulate the semiparametric approach of Ait-Sahalia and Brandt (2001) for determining optimal asset-allocation decisions and show how to extend it to properly account for short-sale constraints. Next, we introduce the different methods for decomposing state variables into their short-term and long-term components.

A. Optimal Portfolio Weights

We consider a pure asset-allocation problem and abstract from intermediate consumption. We endow our single-period investor with a power utility function $U(W_{t+1})=(1-\gamma)^{-1}W_{t+1}^{1-\gamma}$, where W_{t+1} is time $t+1$ wealth, and γ denotes the risk-aversion parameter. The investor picks an asset allocation $x_t \in \mathbb{R}^{1\times m}$ by maximizing expected utility,

$$\text{(1)} \qquad \max_{x_t} \mathrm{E}[U(W_t(R_f + x_t \widetilde{R}_{t+1})) \,|\, Z_t],$$

where R_f denotes the gross risk-free return, $Z_t \in \mathbb{R}^{k\times 1}$ denotes a vector of state variables, and $\widetilde{R}_{t+1} \in \mathbb{R}^{m\times 1}$ denotes the vector of excess returns on the risky assets. For expositional purposes, we assume there is only one risky asset ($m=1$) such that x_t and $\widetilde{R}_{t+1}$ are scalars. We relax this assumption in the empirical analysis in Section IV.

The asset allocation x_t solves the standard first-order condition

$$\text{(2)} \qquad \mathrm{E}[m_{t+1}|Z_t] \quad = \quad 0,$$

where

$$\text{(3)} \qquad m_{t+1} \quad = \quad U'(W_t(R_f + x_t \widetilde{R}_{t+1})) \times W_t \times \widetilde{R}_{t+1}.$$

This implies that x_t is a function of the state variables Z_t. The precise functional form of the relationship between Z_t and x_t is governed by the shape of the utility function and the conditional (on Z_t) return distribution of the risky asset. Rather than assuming a particular parametric form for the latter, we follow Brandt (1999) and Ait-Sahalia and Brandt (2001) and estimate x_t semiparametrically as $x_t = x(\beta Z_t)$, where $\beta \in \mathbb{R}^{1\times k}$ is a parameter vector, and $x: \mathbb{R} \to \mathbb{R}$ is a univariate[5] function that is estimated nonparametrically from the data. The semiparametric approach avoids many potential biases due to misspecification of an assumed functional relationship between Z_t and x_t; see also the discussion in Ait-Sahalia and Brandt (2001).

The first step in estimating the functional form of $x(\cdot)$ is to condition on βZ_t rather than Z_t itself and use sample analogues of the objective function in equation (1) and the first-order condition in equation (2),

$$\text{(4)} \qquad \max_{x(z)} \mathrm{E}\left[U(W_t(R_f + x(z)\widetilde{R}_{t+1})) \,\middle|\, \beta Z_t = z\right]$$

$$\approx \quad \max_{x(z)} \sum_{t=1}^{n} \omega_t(z)\, U(W_t(R_f + x(z)\widetilde{R}_{t+1})),$$

and

$$\text{(5)} \qquad \mathrm{E}[m_{t+1}|\beta Z_t = z] \quad \approx \quad \sum_{t=1}^{n} \omega_t(z)\, m_{t+1}(z) \quad = \quad 0,$$

where n denotes the number of observations, $m_{t+1}(z)$ is defined in equation (3) with x_t replaced by $x(z)$, and $\omega_t(z) = \omega(\beta Z_t, z)$ denotes a weight function. The weights $\omega(\beta Z_t, z)$ depend on the distance of z to βZ_t. If βZ_t is close to z, then this observation contains more information on the optimal conditional asset allocation $x(z)$, and the corresponding weight $\omega_t(z)$ is higher. As a result, the observation receives a higher weight in both equations (4) and (5) and thus has a higher impact on $x(z)$. This is most easily seen if βZ_t can only take one out of two possible values, say, z_1 and z_2, and if $\omega(\beta Z_t, z) = 1$ if $\beta Z_t = z$, and 0 otherwise. In that case, $x(z_1)$ and $x(z_2)$ follow from equation (4), with $x(z_1)$ only based on those observations t for which $\beta Z_t = z_1$ and $x(z_2)$ based on the remaining observations. In our empirical analysis in Section IV, we use a normal kernel for the weight function ω with an adaptive bandwidth choice proposed by Abramson (1982). More details are provided in the Internet Appendix.

To solve for the optimal asset allocation $x(z)$ using either equation (4) or equation (5), we need to estimate the parameter vector β. Not all elements of

[5]If we allow for more risky assets ($m > 1$) as in the empirical application, we have m different functions $x(\cdot)$ that we can estimate nonparametrically.

β are identified due to the nonparametric nature of the asset-allocation function $x(\cdot)$. In particular, we cannot identify any intercept term in βZ_t or the length $\|\beta\|$. Both of these are included in the estimate of $x(\cdot)$. Therefore, we assume that Z_t does not contain a constant term, and we impose the estimation restriction $\|\beta\| = (\beta\beta')^{1/2} = 1$. We can now estimate β using the GMM; see Hansen (1982). Following Ait-Sahalia and Brandt (2001), we use the conditional moment condition in equation (5) to construct the unconditional moment condition

$$\text{(6)} \qquad \mathrm{E}\left[m_{t+1}(\beta Z_t) \otimes g(Z_t)\right] \quad = \quad 0,$$

where $g(Z_t)$ is a known deterministic vector function of Z_t, and $\otimes$ denotes the Kronecker product. In Section IV, we use $g(Z_t) = (1, Z_t')'$. Ait-Sahalia and Brandt (2001) argue that this choice for $g(\cdot)$ is adequate and numerically efficient in the asset-allocation context.

To implement the GMM estimator for β, we approximate equation (6) by

$$\text{(7)} \qquad \bar{m} \quad = \quad \frac{1}{n}\sum_{t=1}^{n} m_{t+1}(\beta Z_t) \otimes g(Z_t) \quad = \quad 0,$$

and use the GMM objective function

$$\text{(8)} \qquad \min_{\|\beta\|=1} \bar{m}' W \bar{m},$$

where W is an appropriate positive definite weighting matrix. Following Hansen (1982), the (infeasible) optimal choice of the weighting matrix is given by

$$W \quad = \quad \mathrm{E}[m_{t+1}m_{t+1}' \otimes g(Z_t)g(Z_t)']^{-1}.$$

We use a feasible multistep approach instead. First, we set $\hat{W}^{(1)} = \mathrm{I}$ and obtain a first-step estimate $\hat{\beta}^{(1)}$ of β. Next, we use $\hat{\beta}^{(1)}$ to update the weighting matrix $\hat{W}^{(1)}$ to

$$\text{(9)} \quad \hat{W}^{(2)} \quad = \quad \left(\frac{1}{n}\sum_{t=1}^{n} m_{t+1}\left(\left(\hat{\beta}^{(1)}\right)' Z_t\right) m_{t+1}\left(\left(\hat{\beta}^{(1)}\right)' Z_t\right)' \otimes g(Z_t)\, g(Z_t)'\right)^{-1}.$$

We use $\hat{W}^{(2)}$ as the new weighting matrix in equation (8) to obtain a new estimate $\hat{\beta}^{(2)}$. The process is repeated until convergence. Usually, three or four iterations suffice. Standard errors are computed based on the converged pair $(\hat{\beta}^{(i)}, \hat{W}^{(i)})$.

B. Imposing Short-Sale Constraints

A plain-vanilla implementation of the estimation procedure as described in Section II.A typically results in optimal asset allocations with unrealistic short positions. This obstructs a sensible interpretation of the results. To avoid such problems, we can impose short-sale constraints on the asset-allocation problem in equation (4). This, however, changes the appropriate moment conditions. In particular, the moment condition in equation (2) no longer holds under short-sale constraints. For example, if an asset allocation lies on the boundary $x_t = 0$, the expectation in the condition in equation (2) is in general strictly negative because further decreasing the asset weight to a short position would increase the objective function. Because the moment conditions in equation (5) also directly enter the GMM criterion of equation (8) via equation (7), we have to account explicitly for the short-sale constraints. Otherwise, our inference of β can be biased.

The correct moment conditions follow from the Kuhn–Tucker first-order conditions for the full Lagrangian of the constrained optimization problem,

$$(10)\quad \max_{0\leq x(z)\leq 1} \mathrm{E}\left[U(W_t(R_f + x(z)\widetilde{R}_{t+1}))\middle|\beta Z_t = z\right] + \lambda_0(z)\times(1-x(z)) + \lambda_1(z)\times x(z),$$

where $\lambda_0(z)$ and $\lambda_1(z)$ are the Lagrange multipliers. If the optimal $x(z)$ provides an internal solution, we have $\lambda_0(z)=\lambda_1(z)=0$, and we recover the original moment condition from Section A. If, however, the constrained optimum lies on the boundary $x(z)=0$, then the Kuhn–Tucker first-order conditions yield $\lambda_0(z)=0$ and $\lambda_1(z)>0$. Similarly, if the optimum lies on the boundary $x(z)=1$, we obtain $\lambda_1(z)=0$ and $\lambda_0(z)>0$. Omitting the Lagrange multipliers from the Kuhn–Tucker first-order conditions therefore gives the wrong moment conditions to identify β. The solution is obvious. Rather than only computing $x(z)$ from the first-order condition, we also compute the Lagrange multipliers $\lambda_i(z)$ for $i=0,\ldots,m$ from equation (10), where m denotes the number of assets. We then use the Lagrange multipliers to redefine the moment condition $\bar{m}$ for the GMM objective function in equation (7) as

$$(11)\qquad \bar{m} = \frac{1}{n}\sum_{t=1}^{n}\tilde{m}_{t+1}(\beta Z_t)\otimes g(Z_t),$$

with

$$(12)\quad \tilde{m}_{t+1}(\beta Z_t) = U'(W_t(R_f + x(\beta Z_t)\widetilde{R}_{t+1}))\times W_t\times\widetilde{R}_{t+1} + \lambda(\beta Z_t) - \lambda_0(\beta Z_t)\iota,$$

where $\lambda(z)=(\lambda_1(z),\ldots,\lambda_m(z))'$ with $\lambda_i(z)\in\mathbb{R}^+\cup\{0\}$ for $i=0,\ldots,m$, and where ι is an $m\times 1$ vector of ones. We provide further details on the numerical implementation in the Internet Appendix.

To illustrate the effect of omitting the Lagrange–Kuhn–Tucker multipliers (LKTMs) from the first-order conditions, we provide an example based on simulated data. In line with our empirical analysis in Section IV, we use the estimated long-term and short-term components of the dividend yield as our two state variables, such that Z_t has two elements. We set $\beta=\beta(\phi)=(\cos(\phi),\sin(\phi))$ for $\phi\in[0,\pi]$, such that $\|\beta\|=1$ by design. We simulate the risky asset returns for stocks and bonds from a standard normal distribution where means and (co)variances are taken from the data in Section IV. We substitute the expected return on stocks by a linear function of βZ_t with $\phi=\pi/3\approx 1$, such that $\beta=(\cos(\phi),\sin(\phi))\approx(0.50,0.87)$. Using the methodology of Section II.A, we use the simulated data to estimate the parameter ϕ and investigate how close it is to the true value $\pi/3$ with and without including the Kuhn–Tucker extensions in the GMM moment conditions.

Figure 1 shows the (scaled) negative value of the GMM criterion function, $-\bar{m}'W\bar{m}$, as a function of ϕ. The two curves correspond to the GMM with and

FIGURE 1

Effect of Short-Sale Restrictions on the GMM Criterion for Simulated Data

Figure 1 presents plots of the (scaled) negative value of the generalized method of moments (GMM) objective function (vertical axis) as a function of the angle ϕ (horizontal axis) with (solid) and without (dashed) inclusion of the Kuhn–Tucker multipliers (KTMs) in the GMM moment conditions. The results are presented for the long-term and short-term components of dividend yield (DY) using the simulation setup described in Section II.A.

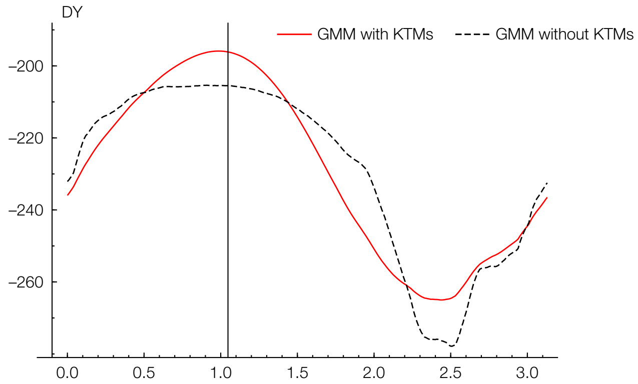

without the inclusion of the LKTMs. The maximum of both curves lies near the true value $\phi \approx 1$, namely, $\hat{\phi} = 0.99$ if the multipliers are included and $\hat{\phi} = 0.91$ if the multipliers are excluded. The effect of the multipliers on the point estimates thus appears to be limited. The curvature of the two objective functions, however, is markedly different. This has a substantial effect on the standard errors. If the multipliers are included, the standard error of $\hat{\phi}$ is 0.015, whereas it is 1.138 if the multipliers are excluded. The difference in magnitude in this simulated example is due to the strength of the relationship between βZ_t and the expected returns on the risky asset. For empirical data, the effect is less drastic but still substantial, with variances increasing by a factor of 2 to 5 if the multipliers are (incorrectly) omitted during the estimation stage. A correct specification of the moment conditions that accounts for short-sale constraints thus has a major impact on the significance of alternative state variables. In all of our empirical analyses, we therefore use the augmented moment conditions based on equations (11) and (12).

C. Filters for Short-Term and Long-Term Components

Decomposing a time series y_t into a long-term or trend component τ_t and a short-term or cyclical component c_t has a long history in economics, particularly in macroeconomics. Familiar techniques include the Christiano–Fitzgerald (CF) filter and the earlier Hodrick–Prescott (HP) filter, as well as alternatives based on unobserved-components time-series models, as explained by Harvey (1990). The benchmark for our analysis is the CF filter (see Christiano and Fitzgerald (2003) and Baxter and King (1999)). The CF filter is a bandpass filter that formulates the trend-cycle decomposition $y_t = \tau_t + c_t$ in the frequency domain. In the CF filtering approach, for each observation time t, the trend (τ_t) and cyclical (c_t) components are essentially a weighted moving average of all observations, backward and forward. This makes the filter actually a *smoother*. Only at the sample end points does the CF filter truly behave as a filter. In our empirical analysis, this renders the decomposition prone to a look-ahead bias. To avoid this, we run the filter recursively after a burn-in period of 9 years and (again recursively) compute the end-of-sample decomposition $\hat{\tau}_t$ and $\hat{c}_t$ for each monthly observation time t. Because the CF filter is optimal for an underlying random-walk time series y_t, this assigns relatively more weight to the end-of-sample observations.

The main advantage of the CF filter is that it allows us to be explicit about the definition of the cyclical component c_t. More precisely, we can define c_t to capture the cyclical variation in y_t up to p months, where p can be chosen by the user. In Section IV, we use a range of values for p from 12 to 48 months to test the sensitivity of our results to the definition of the short-term component. We use the notation CF(p) to denote a CF filter with a short-term component of p months.

In our robustness analysis in Section VI.A, we also consider the results for the alternative of an HP filter; see Leser (1961) and Hodrick and Prescott (1997). This filter is widely used in the context of macro data for trend-cycle decompositions. Like the CF filter, the HP filter also is actually a smoother and would be prone to a look-ahead-bias in an investment context. Therefore, we use the same recursive implementation of the HP filter as for the CF filter. The HP filter has a key tuning parameter λ that determines the amount of smoothness that is imposed on the trend and cycle components. Familiar values are $\lambda = 1{,}600$ for quarterly and $\lambda = 14{,}400$ for monthly data; see Hodrick and Prescott (1997) and Backus and Kehoe (1992). In order not to be sensitive to a particular value of λ, we investigate our results for the HP filter with a range of different values for λ, namely, 100, 1,600, and 14,400.

III. Data

We consider three assets: stocks, bonds, and Treasury bills. Stock returns are given by the monthly returns on the Center for Research in Security Prices (CRSP) value-weighted index for all New York Stock Exchange (NYSE), American Stock Exchange (AMEX), and National Association of Securities Dealers Automated Quotations (NASDAQ) stocks. For bond returns, we use a monthly time series on 10-year Treasury constant-maturity rates from the Federal Reserve Board to approximate the returns on 10-year government bonds. We construct the return at time $t+1$ by calculating the time $t+1$ present value of a 10-year bond issued at par at time t with coupon r_t. The discount rate at time $t+1$ equals the 10-year constant-maturity rate r_{t+1}. For the risk-free rate, we take the 1-month Treasury bill rate. The returns are sampled at a monthly frequency from Jan. 1954 to Dec. 2012. The complete sample has 708 observations.

We consider six popular[6] state variables from the literature: the default spread, the natural log of the dividend yield of the Standard & Poor's (S&P) index, the term spread, a trend or momentum variable for the S&P index, the short rate, and the dividend growth rate. The default spread is measured as the yield difference between Moody's Baa- and Aaa-rated corporate bonds. The dividend yield

[6]Similar variables were used by, for example, Fama and French (1988), (1989), Keim and Stambaugh (1986), Kandel and Stambaugh (1996), Barberis (2000), Campbell and Viceira (1999), (2003), Ait-Sahalia and Brandt (2001), Brandt and Santa-Clara (2006), Handa (2006), and Wachter and Warusawitharana (2009).

is the sum of dividends paid on the S&P index over the past 12 months divided by the current level of the index. The term spread is the yield difference between 10-year and 1-year government bonds. The trend variable is the difference between the natural log of the current S&P index level and the natural log of the average index level over the previous 12 months.

Table 1 provides descriptive statistics for the data. Stock and bond returns both appear to be moderately skewed, but to opposite sides. As always, pairwise correlations of the risky assets' returns with the candidate state variables are modest: The predictable component of returns is typically small compared with return volatility.

TABLE 1
Descriptive Statistics

Table 1 shows descriptive statistics for monthly returns on the Center for Research in Security Prices (CRSP) stock index, a portfolio of 10-year government bonds, and the 1-month Treasury bill rate. The table also includes descriptive statistics for six candidate state variables: the natural log of the dividend yield of the Standard & Poor's (S&P) 500 index (DY); the stock market trend (TR), defined as the difference of the natural log of the index level and its 12-month moving average; the term spread (TS), defined as the yield difference between 10-year and 1-year government bonds; the default spread (DS), defined as the difference between the Baa and Aaa yield; the short rate (SR); and the (annual) dividend growth rate (DG). The data are sampled monthly from Jan. 1963 to Dec. 2012.

	Variables								
Statistic	Stocks	Bonds	T-Bills	DY	TR	TS	DS	SR	DG
Mean	0.89	0.59	0.42	1.03	3.18	0.98	1.04	5.17	0.45
Median	1.22	0.50	0.41	1.10	4.86	0.89	0.92	5.00	0.50
Std. Dev.	4.49	1.99	0.25	0.41	9.87	1.18	0.47	3.04	0.57
Skew	−0.51	0.60	0.68	−0.34	−1.20	−0.07	1.62	0.68	−1.48
Kurtosis	1.93	3.54	1.21	−0.69	2.68	−0.36	3.67	1.18	5.42
Minimum	−22.64	−8.16	0.00	0.10	−46.41	−3.07	0.32	0.01	−2.33
Maximum	16.61	10.74	1.35	1.83	24.52	3.40	3.38	16.30	1.85
Correlation with:									
Bonds	−0.06								
T-bills	−0.02	0.11							
DY	0.09	0.05	0.61						
TR	0.01	−0.08	0.03	−0.13					
TS	0.07	0.05	−0.63	−0.23	0.02				
DS	0.06	0.11	0.23	0.39	−0.28	0.19			
SR	−0.02	0.09	0.98	0.63	0.04	−0.65	0.23		
DG	−0.05	−0.01	0.15	0.07	0.13	−0.26	−0.22	0.15	

IV. Estimation Results

To investigate the relative importance of short-term and long-term components of state variables for asset allocation, we proceed as follows: First, we decompose each of the 6 state variables introduced in Section III into their short-term (c_t) and long-term (τ_t) components using the recursive implementation of the CF filter described in Section II.C. Next, we stack the two components into a bivariate vector of state variables Z_t. Finally, we parameterize the parameter vector $\beta \in \mathbb{R}^{1\times 2}$ as

$$\beta = (\cos(\phi), \sin(\phi))', \qquad 0 \leq \phi < \pi, \tag{13}$$

such that $\|\beta\| = 1$ for every ϕ. We then use the GMM criterion to find the optimal ϕ. Note that the GMM objective function is periodic in ϕ with period π, as can be seen in Figure 1. This is a result of the fact that the GMM criterion is the same for β and $-\beta$ due to the nonparametric estimation of the asset-allocation function $x(\cdot)$.

The time-series standard deviation σ_τ of the long-term component τ_t is typically much larger than that of the short-term component c_t (i.e., than σ_c). This has a direct impact on the magnitude of the elements of β and therefore on ϕ. Because the GMM criterion is already highly nonlinear in the parameter ϕ, the difference in scale adds further complexity to the estimation problem. To avoid this issue and enhance the numerical stability of the estimation algorithm, we scale the short-term and long-term components in Z_t by their standard deviations as

$$(14) \qquad Z_t = \left(\frac{\tau_t - \mu_\tau}{\sigma_\tau}, \frac{c_t - \mu_c}{\sigma_c}\right)',$$

where μ_τ and μ_c are the time-series means of τ_t and c_t over the entire in-sample estimation period, respectively.[7] Note that for the special case $\phi^{\text{ND}} = \arctan(\sigma_c/\sigma_\tau)$, we recover the nondecomposed (ND) original state variable. We can thus test whether the effect of the decomposition is statistically significant by testing $H_0: \phi = \phi^{\text{ND}}$.

We present the estimation results for each state variable and for three versions of the CF filter in Table 2. For the CF(12) filter, the short-term component

TABLE 2

Optimal Component Weights for the Christiano–Fitzgerald Filter

Table 2 presents the generalized method of moments (GMM) estimates of the angle ϕ and weights $\beta = (\beta_\tau, \beta_c) = (\cos(\phi), \sin(\phi))$ for the long-term (τ_t) and short-term (c_t) components of a single state variable. The state variables considered are the natural log of the dividend yield (DY), stock market trend (TR), term spread (TS), default spread (DS), short rate (SR), and dividend growth rate (DG). Standard errors of $\hat{\phi}$ are provided in parentheses. We use a recursive implementation of the Christiano–Fitzgerald (CF) filter for decomposing each state variable, which is denoted as CF(p), where the cyclical (short-term) component is defined as the component with a cycle up to p months for $p = 12, 24, 48$.

Parameter	DY	TR	TS	DS	SR	DG
Panel A. CF(12)						
$\hat{\phi}$	1.951	1.221	0.370	1.621	0.870	2.911
	(0.271)	(0.334)	(0.352)	(0.818)	(0.339)	(0.445)
β_τ	−0.371	0.343	0.932	−0.050	0.644	−0.974
β_c	0.929	0.939	0.362	0.999	0.765	0.228
Panel B. CF(24)						
$\hat{\phi}$	2.231	1.091	0.130	2.161	1.161	2.791
	(0.451)	(0.344)	(0.226)	(0.679)	(0.322)	(0.513)
β_τ	−0.613	0.462	0.992	−0.557	0.399	−0.939
β_c	0.790	0.887	0.130	0.831	0.917	0.343
Panel C. CF(48)						
$\hat{\phi}$	2.321	1.541	0.230	2.601	0.840	3.082
	(0.434)	(0.402)	(0.290)	(0.388)	(0.568)	(0.431)
β_τ	−0.682	0.030	0.974	−0.858	0.667	−0.998
β_c	0.731	1.000	0.228	0.514	0.745	0.060

[7]This recentering and rescaling only enhances the numerical stability of the estimation approach. Any similarly sized numbers μ_τ, μ_c, σ_τ, and σ_c would result in the same out-of-sample investment performance later on, such that no look-ahead bias is introduced in this way.

https://doi.org/10.1017/S002210901700059X Published online by Cambridge University Press

(β_c) receives more weight than the long-term component (β_τ) for the dividend yield (DY), stock market trend (TR), and default spread (DS). For the DY, DS, and dividend growth (DG), the loadings of the long-term and short-term components have opposite signs. This excludes the possibility that the nondecomposed state variable is optimal for asset allocation, which would require two strictly positive weights. Economically, this implies that investors should react differently to persistent changes in state variables compared with transitory changes. Long-term fundamental changes may reflect secular trends in expected returns and therefore have an impact on optimal asset allocations. Short-term deviations, however, can be the consequence of predictability in excess of expected returns due to, for instance, overreaction of markets to short-term information. As a result, such information should be exploited differently; see also Koijen et al. (2009).

We present the information contained in Table 2 graphically in Figure 2, where we plot $\hat{\phi}$ and its 90% confidence band as a function of $p=12,24,48$. Each chart corresponds to a different state variable. The horizontal lines in the figure reflect different null hypotheses of interest, namely: i) Only the long-term component matters (H_0: $\phi=0$ or $\phi=\pi$); ii) only the short-term component matters (H_0: $\phi=\pi/2$); iii) the two scaled components are weighted equally (H_0: $\phi=\pi/4$); and iv) the nondecomposed state variable is the relevant state variable (H_0: $\phi=\phi^{\mathrm{ND}}$).

The pattern for the natural log of the dividend yield (DY) in Figure 2 shows that the value of $\hat{\phi}$ changes gradually if we extend the definition of the short-term component from 1 ($p=12$) to 2 ($p=24$) to 4 ($p=48$) years. For $p=12$ and $p=24$, we cannot reject the hypothesis that only the short-term component matters (i.e., H_0: $\phi=\pi/2$); the horizontal line $\phi=\pi/2$ lies inside the confidence band. However, for all filters, we can reject the hypothesis that only the long-term component is relevant ($\phi=0$ or $\phi=\pi$); both $\phi=0$ and $\phi=\pi$ always lie outside the 90% confidence band. For the same reason, we can reject that an unweighted average of the short-term and long-term components is optimal ($\phi=\pi/4$) and that the nondecomposed state variable is optimal ($\phi=\phi^{\mathrm{ND}}$).

A similar result holds for the stock market trend (TR). For all filters, we can reject the hypothesis that only the long-term component matters for optimal portfolio choice, whereas we cannot reject that only the short-term component matters. Also note that we can reject the optimality of the nondecomposed state variable ($\phi=\phi^{\mathrm{ND}}$) for all filters considered. The decomposition of the stock market trend into its short-term and long-term components is thus actually statistically significant. The results for the short rate (SR) are similar. However, for $p=48$ (4 years), the decomposed state variable no longer performs significantly better than its nondecomposed counterpart. By contrast, for the CF(12) filter, we see that neither the long-term nor the short-term component is sufficient and that we need both components in the optimal index βZ_t.

The importance of the short-term components in the dividend yield, stock market trend, and short rate stand in sharp contrast to the results for the term spread (TS) and dividend growth (DG). For both TS and DG, we cannot reject that only the long-term component matters, whereas we can reject that only the short-term component matters. In addition, for the term spread and the dividend growth rate, we cannot reject the hypothesis that the nondecomposed state variable is

FIGURE 2

Optimal Decomposition Angles for Alternative Christiano–Fitzgerald Filters

Figure 2 presents the optimal angle $\hat{\phi}$ (solid with diamond) and its 90% confidence band (dashed) for the long-term and short-term component weights $\beta_\tau = \cos(\phi)$ and $\beta_c = \sin(\phi)$, respectively, for different state variables and filters. The horizontal axis gives the value of p, where CF(p) is the recursive Christiano–Fitzgerald filter used for the GMM estimation of ϕ. The graphs are for the different state variables: natural log of the dividend yield (DY), stock market trend (TR), term spread (TS), default spread (DS), short rate (SR), and dividend growth rate (DG). The horizontal lines in each chart correspond to the different null hypotheses of interest: i) Only the long-term component matters (H_0: $\phi = 0$ or $\phi = \pi$); ii) only the short-term component matters (H_0: $\phi = \pi/2$); iii) the two scaled components are weighted equally (H_0: $\phi = \pi/4$); and iv) the nondecomposed state variable is the relevant state variable (H_0: $\phi = \phi^{\mathrm{ND}}$, with $\phi^{\mathrm{ND}} = \arctan(\sigma_\tau/\sigma_c)$).

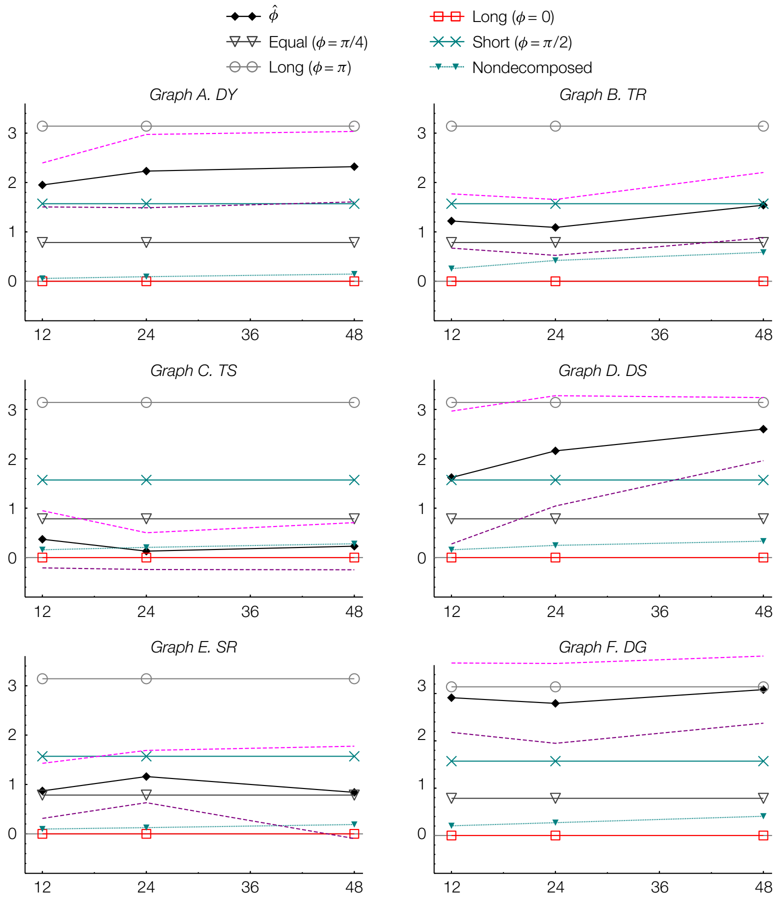

adequate.[8] Finally, the results for the default spread (DS) vary with p. For $p = 12$, the short-term component ($\phi = \pi/2$) lies inside the confidence band. For $p = 48$, in contrast, the long-term component lies in the confidence band, whereas the short-term component does not.

[8]This may be less clear at first sight for dividend growth. Due to the periodicity of the GMM criterion in ϕ, however, we can shift the curve for ϕ^{ND} up by π, which causes it to lie inside the confidence band.

To summarize, the decomposition into short-term and long-term components turns out to be relevant for asset allocation for a number of familiar state variables from the literature. The results are most robust for the dividend yield and stock market trend, followed by the short rate. The evidence in favor of the relevance of short-term components is also most strong for these three variables. The relevance of long-term components and of nondecomposed state variables, conversely, is most pressing for the term spread and dividend growth rate. These variables are also more likely related to secular trends in risk premia rather than to the short-term predictability of returns.

V. Performance of Investment Strategies

Using the optimal weight vectors $\hat{\beta}$ as estimated in Section IV, we compute the optimal asset allocations $x(\hat{\beta}Z_t)$ as a function of the index $\hat{\beta}Z_t$. The results are presented in Figure 3 for the CF(12) filter. To facilitate a visual comparison between state variables, we standardize each βZ_t on the horizontal axis by its time-series standard deviation.

The optimal asset allocations for each state variable clearly depend on the value of $\hat{\beta}Z_t$ (i.e., the linear combination of long-term and short-term components). For some state variables, such as the short rate (SR) or dividend growth

FIGURE 3

Optimal Asset Allocations as a Function of the Index βZ_t

Figure 3 presents the optimal asset allocation as a function of the index $\hat{\beta}Z_t$ computed using the optimal $\hat{\beta}$ from Table 2. The horizontal axis corresponds to $\hat{\beta}Z_t$, standardized by its time-series standard deviation. The CF(12) filter is used for the decomposition. The graphs are for the different state variables: natural log of the dividend yield (DY), stock market trend (TR), term spread (TS), default spread (DS), short rate (SR), and dividend growth rate (DG).

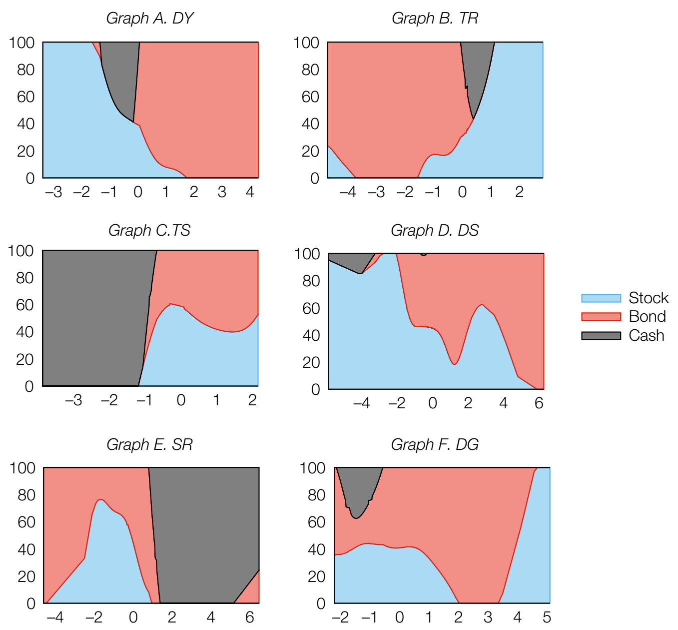

(DG), the patterns are highly nonlinear. For other variables, the pattern is more straightforward. For the term spread (TS), for example, the optimal allocation switches from a pure cash strategy to an equal mix of stocks and bonds for an increasing index βZ_t. For the dividend yield (DY) and default spread (DS), the allocation switches from stocks for low values of the index, via stocks and cash for medium values of the index, to bonds only for high values of the index. The reverse pattern holds for the stock market trend.

To understand these results, we note that the short-term component receives the dominant weight in the asset-allocation decision for both the dividend yield and stock market trend. As stated in the Introduction, the short-term component of the dividend yield or stock market trend correlates with short-term momentum effects. For example, given the persistence of dividends, a large temporary negative innovation in the dividend yield as picked up by the short-term component corresponds to a positive shock to the stock price. From a momentum-based perspective, we would then invest a large portion of our wealth in stocks. This is precisely what happens due to the large weight of the short-term component: Negative values of the short-term component correlate with low values of index βZ_t and with a 100% investment in stocks. The converse holds for a large positive innovation in the dividend yield, as picked up by the short-term component. The momentum interpretation and the positive loading for the short-term component are therefore in line with the results of Koijen et al. (2009) for a short-term (1-month) investor. The long-term component of the dividend yield, by contrast, picks up the secular trends in the dividend yield over time and mainly corresponds to the mean-reversion effect. It receives a corresponding negative weight β_τ in the optimal asset-allocation decision.

The asset allocations from Figure 3 can also be implemented in a backtest. At the start of each month t we use the asset allocation $x(\hat{\beta} Z_t)$ and hold it over the next month. To prevent a look-ahead bias due to estimation, we estimate the model over the period Jan. 1963–Dec. 2002 and implement the induced asset-allocation strategies out-of-sample for each month over the period Jan. 2003–Dec. 2012 while keeping $\hat{\beta}$ fixed. For each combination of a filter and state variable, we compute monthly average returns, standard deviations, realized Sharpe ratios, and average utilities. We present the results in Table 3.

The top entries in Table 3 present the results for the original, nondecomposed state variables. These results serve as our benchmark. We also consider three alternative benchmark portfolios corresponding to 100% stocks, 100% bonds, and 50%/50% stocks and bonds. The Sharpe ratios of these benchmarks are 0.135, 0.166, and 0.209, respectively. Comparing these with the Sharpe ratios for the nondecomposed state variables in the top panel, we see that most of the state variables have a very limited explanatory power out-of-sample. Only for the stock market trend (TR) do we observe improvements in the Sharpe ratios in comparison with the naive strategy of 50%/50% stocks and bonds.

If we decompose the state variables into their short-term and long-term components, we see that the results improve markedly for the dividend yield (DY) and the stock market trend (TR) compared with using the nondecomposed state variables. Given the investor's 1-month holding period, the best results are achieved with the CF(12) and CF(24) filters, which define the short-term component to be

TABLE 3
Investment Performance for the Christiano–Fitzgerald Filter

Table 3 presents the out-of-sample performance of different portfolio strategies using the recursively implemented Christiano–Fitzgerald (CF) filter for $p=12,24,48$. The state variables are the natural log of the dividend yield (DY), stock market trend (TR), term spread (TS), default spread (DS), short rate (SR), and dividend growth rate (DG). The model is estimated over the period Jan. 1963–Dec. 2002 and evaluated over the period Jan. 2003–Dec. 2012 using a holding period of 1 month. We report average returns ($\bar{R}$), standard deviations (σ), Sharpe ratios (SHARPE), and average power utilities ($\bar{U}$) with the risk-aversion parameter $\gamma=5$. For each state variable, we test the hypothesis that the average return, standard deviation, Sharpe ratio, or average utility is significantly higher for the decomposed versus nondecomposed state variables. * indicates significance at the 5% level. For testing the significance of Sharpe ratio improvements, the approach of Jobson and Korkie (1981) is used with the correction suggested by Memmel (2003).

Statistic	DY	TR	TS	DS	SR	DG
Panel A. Nondecomposed						
$\bar{R}$	0.004	0.005	0.003	0.002	0.004	0.002
σ_R	0.018	0.020	0.022	0.029	0.021	0.019
SHARPE	0.207	0.262	0.117	0.066	0.194	0.089
$\hat{U}$	0.0043	0.0055	0.0027	0.0008	0.0043	0.0021
Panel B. CF(12)						
$\bar{R}$	0.006	0.007	0.003	0.002	0.004	0.001
σ_R	0.019	0.019	0.022	0.026	0.035	0.019
SHARPE	0.326	0.342	0.134	0.058	0.124	0.041
$\hat{U}$	0.0066*	0.0070	0.0031	0.0008	0.0023	0.0011
Panel C. CF(24)						
$\bar{R}$	0.005	0.006*	0.003	0.001	0.004	0.002
σ_R	0.017	0.021	0.021	0.030	0.038	0.018
SHARPE	0.294	0.271	0.125	0.036	0.112	0.105
$\hat{U}$	0.0056	0.0059	0.0029	−0.0002	0.0017	0.0024
Panel D. CF(48)						
$\bar{R}$	0.005	0.005	0.002	0.001	0.003	0.002
σ_R	0.021	0.020	0.021	0.029	0.030	0.018
SHARPE	0.247	0.265	0.104	0.035	0.112	0.114
$\hat{U}$	0.0055	0.0056	0.0024	−0.0001	0.0023	0.0026

up to 1 and 2 years, respectively. The Sharpe ratios for the dividend yield (DY) and stock market trend (TR) increase by 57% and 30% compared with the nondecomposed benchmark for decompositions with a "short-term" component of up to 1 year, whereas the average expected utilities $\hat{U}$ improve by 53% and 27%, respectively. For all state variables, the advantage of using the state-variable decomposition tapers off if the definition of the short-term component is extended from 1 or 2 years to 4 years. This is in line with our expectations given the 1-month holding period: Valuable signals are more likely to be attributable to components with a stricter short-term definition.

We also test formally whether the performance measures (mean, standard deviation, Sharpe ratio, and average utility) for decomposed state variables are significantly higher than those for the nondecomposed state variables. For the dividend yield, the average return of the decomposed strategy using the CF(12) filter is higher, but the difference is insignificant at the 5% level. The improvement in the average return also leads to an (insignificant) improvement in the Sharpe ratio[9] and a statistically significant (at 5%) rise in the average utility. We obtain similar, although slightly weaker, results for the stock market trend. The other state variables reveal no clear out-of-sample improvements in investment performance.

[9]The differences in average returns and Sharpe ratios are only significant at the 10% level.

Using the CF(24) filter also leads to higher Sharpe ratios and average utilities for the dividend yield and stock market trend, although the increases are insignificant at the 5% level. Again, the improvements appear to be mainly due to increases in average returns rather than to reductions in risk. The increase in the average return is statistically significant for the stock market trend. When we extend the definition of the short-term component to 4 years when using the CF(48) filter, we no longer obtain clear improvements.

The results for dividend yields and the stock market trend are in line with earlier findings. The dividend yield is one of the most widely mentioned candidates when discussing stock return predictability, and the stock market trend can be interpreted as a proxy for the momentum effect. Both the dividend yield and the stock market trend thus directly reflect information related to stock market developments. As stated in the Introduction, the short-term components of these two state variables can be interpreted as a proxy for momentum, which is a dominant factor in portfolio choice for a short-term investor; see Koijen et al. (2009). Therefore, the short-term components of these state variables receive a larger weight in asset-allocation decisions. The decomposition of these state variables into their persistent and transitory components enables investors to exploit the short-term predictability more efficiently and makes the decomposition valuable.

VI. Robustness Checks

A. Alternative Decomposition Filters

The results in Sections IV and V are based on the CF filter as the more recent bandpass filter. A well-known and widely used alternative filter is that of Hodrick and Prescott (1997). We repeat some of our earlier analyses for this alternative filter to check the robustness of our results. We implement the HP filter recursively using smoothing parameter values $\lambda = 100, 1{,}600, 14{,}400$, corresponding to the optimal choice for annual, quarterly, and monthly data, respectively.

The estimation results for the optimal index weights β (not shown) are very similar to those of the CF filter. For the dividend yield, stock market trend, and short rate, the weight β_c of the short-term component is larger in absolute size than that of the long-term component. For the term spread and dividend growth rate, the opposite holds.

We present the out-of-sample performance of the investment strategies based on the HP filter in Table 4. The results are highly robust and comparable to the out-of-sample performance reported for the CF filter. For the dividend yield and stock market trend, average returns, Sharpe ratios, and average utilities increase when using the decomposed rather than the nondecomposed state variables. For example, when using the dividend yield as a state variable and $\lambda = 100$ as the smoothing parameter, the average return, Sharpe ratio, and expected utility increase by 42%, 31%, and 34%, respectively, although none of these differences is statistically significant at the 5% level.[10] Similar but slightly stronger results hold for the stock market trend. For the state variables that do not contain direct information related

[10] Again the differences in average returns and average utilities are only significant at the 10% level.

TABLE 4
Investment Performance for the Hodrick–Prescott Filter

Table 4 is the analogue of Table 3 but for the Hodrick–Prescott (HP) filter. See the explanatory note to Table 3. The tuning constant λ of the HP(λ) filter is set to 100, 1,600, and 14,400, corresponding to the optimal choice for annual, quarterly, and monthly data, respectively.

Statistic	DY	TR	TS	DS	SR	DG
Panel A. Nondecomposed						
$\bar{R}$	0.004	0.005	0.003	0.002	0.004	0.002
σ_R	0.018	0.020	0.022	0.029	0.021	0.019
SHARPE	0.207	0.262	0.117	0.066	0.194	0.089
$\hat{U}$	0.0043	0.0055	0.0027	0.0008	0.0043	0.0021
Panel B. HP(100)						
$\bar{R}$	0.005	0.007	0.003	0.001	0.005	0.001
σ_R	0.020	0.019	0.023	0.028	0.039	0.019
SHARPE	0.272	0.381	0.113	0.041	0.132	0.067
$\hat{U}$	0.0058	0.0076	0.0026	0.0003	0.0024	0.0016
Panel C. HP(1,600)						
$\bar{R}$	0.006	0.006	0.002	0.001	0.004	0.002
σ_R	0.021	0.019	0.021	0.031	0.038	0.018
SHARPE	0.283	0.325	0.115	0.036	0.097	0.108
$\hat{U}$	0.0061	0.0065	0.0027	−0.0002	0.0010	0.0024
Panel D. HP(14,400)						
$\bar{R}$	0.006	0.006	0.002	0.002	0.001	0.002
σ_R	0.022	0.022	0.022	0.032	0.031	0.019
SHARPE	0.287	0.260	0.104	0.048	0.039	0.105
$\hat{U}$	0.0064	0.0057	0.0024	0.0000	−0.0002	0.0024

to the stock market, the decomposition again does not contribute to a stronger out-of-sample investment performance.

B. Different Levels of Risk Aversion

In this section, we check the effect of different risk appetites of investors on the optimal investment strategies and on the optimal relative weights of the long-term and short-term components.

So far, we considered an investor with a constant relative risk aversion (CRRA) utility function with the relative risk-aversion parameter $\gamma = 5$. Here we also consider levels of $\gamma = 1, 5, 10$.

The Internet Appendix contains the estimates of the optimal index weights $\hat{\beta}$ and the corresponding optimal asset allocations. The index weights are very similar to those for $\gamma = 5$ and therefore are not reported here. If anything, the short-term component becomes slightly less important if we increase γ. This holds for all state variables, except for dividend growth. The optimal asset allocations also reveal no major surprises. For example, for the dividend yield and $\gamma = 1$, the optimal asset allocation is 100% stocks for low values of the index and 100% bonds for high values of the index. The switch from one to the other is quite abrupt. As the level of risk aversion increases, the switch becomes more gradual, and we also invest more in cash and bonds for the middle and high index values.

We present the investment performance in Table 5. For the dividend yield and stock market trend and investors with different levels of risk aversion, we observe almost the same pattern as before. Investment strategies based on decomposed state variables outperform those based on their nondecomposed counterparts.

TABLE 5
Investment Performance for Different Risk-Aversion Levels

Table 5 is the analogue of Table 3 and presents the out-of-sample performance of portfolio strategies. The table considers the results for constant relative risk-aversion (CRRA) utility functions with the relative risk-aversion parameter $\gamma = 1, 5, 10$. The CF(12) filter is used for the decomposition. The weights used for βZ_t are those of the recursive decomposition from Table A1. The state variables are the natural log of the dividend yield (DY), stock market trend (TR), term spread (TS), default spread (DS), short rate (SR), and dividend growth rate (DG).

Statistic	DY	TR	TS	DS	SR	DG
Panel A. $\gamma = 1$						
Nondecomposed						
$\bar{R}$	0.005	0.006	0.002	0.005	0.006	0.004
σ_R	0.042	0.033	0.038	0.041	0.041	0.041
SHARPE	0.123	0.189	0.062	0.124	0.140	0.101
$\hat{U}$	0.0057	0.0070	0.0030	0.0056	0.0062	0.0046
CF(12)						
$\bar{R}$	0.008	0.008	0.005	0.005	0.006	0.002
σ_R	0.029*	0.032	0.044	0.039	0.044	0.036
SHARPE	0.293*	0.260	0.121	0.126	0.136	0.048
$\hat{U}$	0.0093	0.0092	0.0057	0.0055	0.0063	0.0024
Panel B. $\gamma = 5$						
Nondecomposed						
$\bar{R}$	0.004	0.005	0.003	0.002	0.004	0.002
σ_R	0.018	0.020	0.022	0.029	0.021	0.019
SHARPE	0.207	0.262	0.117	0.066	0.194	0.089
$\hat{U}$	0.0043	0.0055	0.0027	0.0008	0.0043	0.0021
CF(12)						
$\bar{R}$	0.006	0.007	0.003	0.002	0.004	0.001
σ_R	0.019	0.019	0.022	0.026	0.035	0.019
SHARPE	0.326	0.342	0.134	0.058	0.124	0.041
$\hat{U}$	0.0066*	0.0070	0.0031	0.0008	0.0023	0.0011
Panel C. $\gamma = 10$						
Nondecomposed						
$\bar{R}$	0.003	0.003	0.003	0.001	0.004	0.001
σ_R	0.015	0.014	0.018	0.022	0.017	0.013
SHARPE	0.236	0.246	0.163	0.040	0.224	0.092
$\hat{U}$	0.0037	0.0038	0.0027	−0.0010	0.0037	0.0017
CF(12)						
$\bar{R}$	0.005	0.005	0.003	0.001	0.003	0.001
σ_R	0.016	0.014	0.018	0.016*	0.022	0.013
SHARPE	0.324	0.333	0.181	0.074	0.122	0.061
$\hat{U}$	0.0053	0.0051	0.0030	0.0011	0.0012	0.0013

We get stronger results for less-risk-averse investors ($\gamma = 1$); the standard deviation of returns is significantly (at 5%) lower in that case compared with the nondecomposed strategy, which leads to a significantly higher Sharpe ratio. For the dividend yield, the improvements in Sharpe ratios are 138%, 57%, and 37% for investors with the risk-aversion parameters $\gamma = 1$, 5, and 10, respectively. The corresponding values for the stock market trend are 38%, 31%, and 35%, respectively. For the other state variables, the decomposed variables again do not result in clear, consistent patterns and substantial out-of-sample gains compared with the nondecomposed state variables.

C. Longer Investment Horizons

We also investigate the robustness of our results with respect to the holding period. We consider holding periods of 1, 3, and 12 months. This is particularly relevant in our context because the short-term component may become less important for investors with a longer holding period.

We use nonoverlapping data because it is not clear how to interpret risk and risk-corrected performance measures such as the Sharpe ratio for overlapping data. Using nonoverlapping data while extending the holding period from 1 to 12 months quickly reduces the number of observations. This poses obvious challenges to the reliable estimation of model parameters and performance measures. To mitigate this problem, we use all the data available for both the estimation and backtesting stages. This has to be kept in mind during the discussion of the results. We still compute the long-term and short-term components recursively, such that the decomposition itself is only based on data actually available at time t. When extending the holding period, we also change the corresponding risk-free rate. For each holding period of 1, 3, or 12 months, we use the corresponding Treasury bill rate as the appropriate risk-free rate. We have a total of 200 and 50 nonoverlapping investment periods for the 3- and 12-month holding periods, respectively.

Table 6 shows the optimal component weights β_τ and β_c for the CF(12) filter. Except for dividend growth (DG), the signs of β_τ and β_c are constant for all state variables across all holding periods. There are some interesting variations in the relative magnitude of the weights. Most notably, for the dividend yield (DY), the weight β_c of the short-term component is lower for the 3- and 12-month compared to the 1-month holding period. The long-term component receives a correspondingly larger (negative) weight β_τ. We attribute this pattern to the fact that for a long-term investor, mean reversion becomes the dominant factor in asset-allocation decisions; compare Koijen et al. (2009). Because the long-term component of the dividend yield captures the persistent changes in the dividend yield and therefore the mean-reversion effect, this component becomes more important (with negative sign) for investors with longer holding periods.

To observe this pattern more clearly, Figure 4 presents the optimal asset allocations over different holding periods for the dividend yield (DY). For both 3-month and 12-month holding periods, the larger negative weight for the long-term component implies that high levels of dividend yields correspond to low levels of the optimal index βZ_t, which correspond to 100% in stocks. This is consistent with the investor following a mean-reversion strategy. The opposite holds for low levels of dividend yields. In that case, the allocation almost exclusively consists of bonds.

For the stock market trend, it seems that the information content of the short-term component still dominates its long-term counterpart even if the holding period increases. This implies that the stock market trend holds little information about secular trends in expected returns and long-term mean-reversion effects in stock prices.

Table 7 presents the investment performance over the different holding periods. In contrast to Table 3, Table 7 is based on all the data and is therefore more in-sample. For the 1-month holding period, the results are slightly inferior in terms of magnitude to the out-of-sample results in Table 7, but they are still robust in that the largest gains are obtained for the dividend yield (DY) and stock market trend (TR). For the dividend yield, we observe significant improvements in the average return, Sharpe ratio, and average utility. For the stock market trend, we find a significantly lower standard deviation of returns for only the decomposed strategy.

TABLE 6

Optimal Component Weights for Different Holding Periods

Table 6 presents the generalized method of moments (GMM) estimates of the optimal weights for the long-term and short-term components for holding periods of 1, 3, and 12 months (see also the explanatory note of Table 2). The estimation period is Jan. 1963–Dec. 2012. The state variables are the natural log of the dividend yield (DY), stock market trend (TR), term spread (TS), default spread (DS), short rate (SR), and dividend growth rate (DG).

	Filter: CF(12)					
Weight	DY	TR	TS	DS	SR	DG
Panel A. 1-Month Holding Period						
β_τ	−0.371	0.343	0.932	−0.050	0.644	−0.974
β_c	0.929	0.939	0.362	0.999	0.765	0.228
Panel B. 3-Month Holding Period						
β_τ	−0.863	0.248	0.978	−0.932	0.852	−0.718
β_c	0.506	0.969	0.209	0.362	0.523	0.696
Panel C. 12-Month Holding Period						
β_τ	−0.836	0.010	0.964	−0.868	0.660	0.990
β_c	0.548	1.000	0.267	0.497	0.752	0.140

FIGURE 4

Optimal Allocations for or Different Investment Horizons

Figure 4 presents the optimal asset allocation as a function of the index $\hat{\beta}Z_t$ computed using the optimal $\hat{\beta}$ from Table 6 and Z_t holding the long-term and short-term components of the natural log of the dividend yield (DY). The different graphs are for 1-month, 3-month, and 12-month buy-and-hold investment horizons. The CF(12) filter is used for the decomposition. The index βZ_t is standardized by its time-series standard deviation.

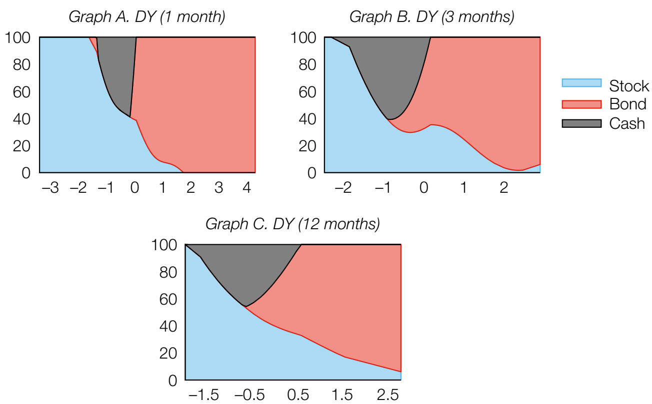

If we move to quarterly (3m) holding periods, a similar pattern emerges. For the stock market trend and dividend yield, the investment strategy based on the decomposed state variables results in Sharpe ratios and average utilities that are (insignificantly) higher than their nondecomposed counterparts. The improvements in Sharpe ratios for both dividend yields and the stock market trend remain robust if we move further out to holding periods of 12 months. Here we observe significantly higher average returns for the dividend yield, which also lead to significantly higher average utility. Again, we see little improvement for the other state variables that are not directly based on stock price information, such as the short rate, term spread, and default spread.

TABLE 7
Investment Performance over Different Holding Periods

Table 7 is the analogue of Table 3 and presents the investment performance for different holding periods. The sample is from Jan. 1963 to Dec. 2012. The state variables are the natural log of the dividend yield (DY), stock market trend (TR), term spread (TS), default spread (DS), short rate (SR), and dividend growth rate (DG).

Statistic	DY	TR	TS	DS	SR	DG
Panel A. 1-Month Holding Period						
Nondecomposed						
$\bar{R}$	0.003	0.005	0.004	0.003	0.004	0.003
σ_R	0.022	0.025	0.022	0.021	0.022	0.020
SHARPE	0.153	0.188	0.200	0.145	0.184	0.154
$\hat{U}$	0.0063	0.0072	0.0073	0.0061	0.0069	0.0062
CF(12)						
$\bar{R}$	0.005*	0.005	0.004	0.004	0.004	0.003
σ_R	0.022	0.022*	0.021	0.023	0.024	0.021
SHARPE	0.218*	0.212	0.198	0.157	0.179	0.147
$\hat{U}$	0.0076*	0.0076	0.0072	0.0065	0.0069	0.0061
Panel B. 3-Month Holding Period						
Nondecomposed						
$\bar{R}$	0.009	0.011	0.013	0.009	0.010	0.009
σ_R	0.035	0.043	0.039	0.040	0.037	0.040
SHARPE	0.267	0.258	0.324	0.225	0.263	0.223
$\hat{U}$	0.0185	0.0188	0.0207	0.0174	0.0184	0.0172
CF(12)						
$\bar{R}$	0.011*	0.013	0.012	0.009	0.011	0.009
σ_R	0.038	0.039	0.039	0.041	0.042	0.038
SHARPE	0.300	0.326	0.322	0.230	0.263	0.240
$\hat{U}$	0.0200	0.0210	0.0206	0.0176	0.0188	0.0178
Panel C. 12-Month Holding Period						
Nondecomposed						
$\bar{R}$	0.035	0.033	0.036	0.034	0.029	0.034
σ_R	0.073	0.079	0.075	0.083	0.071	0.079
SHARPE	0.486	0.425	0.486	0.417	0.405	0.427
$\hat{U}$	0.0651	0.0630	0.0651	0.0631	0.0609	0.0629
CF(12)						
$\bar{R}$	0.046*	0.037	0.035	0.035	0.029	0.033
σ_R	0.081	0.073	0.073	0.081	0.071	0.078
SHARPE	0.565	0.514	0.482	0.426	0.410	0.428
$\hat{U}$	0.0702*	0.0659	0.0647	0.0634	0.0616	0.0629

D. Optimal Filters

So far, we presented the performance of different portfolio strategies using the recursively implemented CF filter for $p=12$, 24, and 48. The choice of p is key to our approach, and robustness of the results to the choice of p is therefore crucial. In this section, we perform the whole out-of-sample recursive investment analysis for all filters $p=3,4,\dots,48$. This is a major computational exercise.

Because we obtained the best out-of-sample results for the dividend yield and stock market trend, we confine our analysis to these two state variables. Figure 5 presents the out-of-sample results for Sharpe ratios and average utilities. Diamonds highlight when the decomposed state variable yields a significantly higher Sharpe ratio or average utility than the nondecomposed state variable. Diamonds are filled or empty depending on whether the difference is statistically significant at the 5% or 10% level, respectively. Similar graphs for other holding periods are provided in the Internet Appendix.

The benchmark Sharpe ratio (or average utility) is that of the nondecomposed approach, which is displayed as the horizontal line in each chart within

FIGURE 5

Sharpe Ratios and Average Utilities of Investment Strategies for All Filters

Figure 5 presents the out-of-sample Sharpe ratios (Graph A) and average utilities (Graph B) of different portfolio strategies using the recursively implemented Christiano–Fitzgerald (CF(p)) filter for $p=3,\ldots,48$ (horizontal axis). State variables are the natural log of the dividend yield (DY, left) and stock market trend (TR, right). The Sharpe ratio and average utility of the nondecomposed strategy are indicated by a horizontal line. Filled and empty diamonds highlight whether the decomposed strategy significantly outperforms the nondecomposed strategy at the 5% and 10% significance levels, respectively.

Graph A. Sharpe Ratios

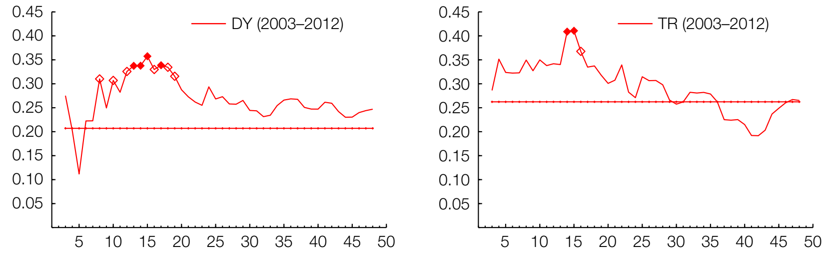

Graph B. Average Utilities

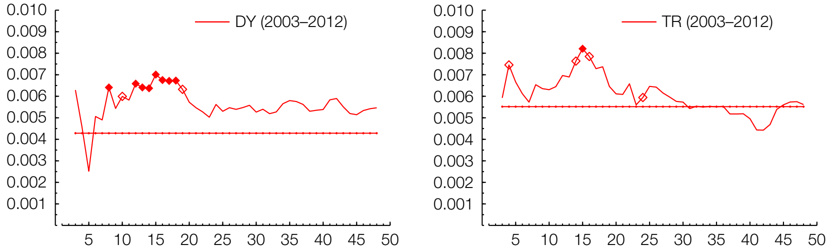

Figure 5. For the dividend yield, we see that the decomposed approach outperforms the nondecomposed approach if we define the short-term component almost anywhere between $p=12$ (12 months) and $p=19$ (19 months). This holds for both the Sharpe ratio and the average utility as the performance metric. The results for the dividend yield are thus robust to the specific choice of $p=12$ in the previous tables, as long as the short-term is defined as a period between 12 and 19 months. For the stock market trend, the decomposed approach outperforms for $p=14$ to $p=16$ using both the Sharpe ratio and the average utility. Similar findings are obtained for holding periods of 3 or 12 months; see the Internet Appendix.

Figure 5 shows that we find significant (at 5%) outperformance (solid diamonds) in terms of both the Sharpe ratio and average utility using roughly $p=13$ to $p=17$ for the dividend yield and using $p=15$ for the stock market trend. In principle, the current out-of-sample Sharpe ratios for different values of p could be used as a selection criterion to determine the "optimal" value of p. Working out the statistical theory of such a selection procedure is daunting due to the different stages of the estimation process and is beyond the scope of the current article. We therefore confine ourselves to presenting the results for the whole range of filters $p=3,\ldots,48$ rather than for one single "optimal" filter.

VII. Conclusion

We decompose a number of commonly used state variables for asset allocation into their long-term and short-term components and investigate whether these components have a different impact on optimal portfolio choice. For an investor with a 1-month investment horizon, we find that the short-term component of state variables that include stock price information receive a larger weight in asset-allocation decisions. This also results in better out-of-sample Sharpe ratios and expected utilities. We attribute the outperformance of the strategies based on the decomposed state variables to short-term market dynamics such as momentum. This is confirmed by the fact that short-term components are less relevant out-of-sample for state variables that do not contain direct price information, such as the short rate, term spread, default spread, or dividend growth rate. Such variables appear to be more connected to secular trends in risk premia.

For investors with longer holding periods, more weight is assigned to the long-term component, particularly for the dividend yield. Because the long-term component in the dividend yield largely picks up secular shifts in expected returns, we conclude that these have a more important role in asset-allocation decisions for longer-term investors. Our approach therefore provides another perspective on the debate about the sources of predictability of financial ratios such as the dividend yield that can be attributed either to irrational mispricing or to rational shifts in underlying risk premia. The results show that investors with a 1-month horizon put more weight on the short-term component (mispricing), whereas investors with a 1-year horizon are relatively more concerned with the long-term component (secular shifts in risk premia).

As a side result, we also make a methodological contribution. We show how short-sale constraints in the asset-allocation problem can influence the formulation of moment conditions in the semiparametric estimation approach of Ait-Sahalia and Brandt (2001). Short-sale constraints require the inclusion of Lagrange–Kuhn–Tucker multipliers in the moment conditions. Omitting these multipliers changes the GMM objective function and can substantially affect the estimates and their standard errors.

Our results are highly robust. Different choices of filters do not alter the results. For all filters considered, the decomposition of the dividend yield and stock market trend turn out to be valuable. The same result is true for changes in the investor's risk-aversion parameter. Interestingly, the results are also confirmed if the holding period of 1 month is extended to 3 or even 12 months. In all of these cases, the Sharpe ratios based on the decomposed dividend yield and stock market trend are higher than those based on their nondecomposed counterparts.

The current results are interesting because they suggest that state variables may contain information at different frequencies with a possibly different relevance for investors. This suggests interesting avenues for further research. In this article, we have made a first step in the decomposition of familiar state variables. Extensions that combine different components of different state variables at the same time can be thought of as a next step. The reliable estimation of all index weights simultaneously for the typical data at hand, however, poses both computational and empirical identification challenges that then have to be tackled first.

References

Abramson, I. S. "On Bandwidth Variation in Kernel Estimates—A Square Root Law." *Annals of Statistics*, 10 (1982), 1217–1223.

Ait-Sahalia, Y., and M. Brandt. "Variable Selection for Portfolio Choice." *Journal of Finance*, 56 (2001), 1297–1351.

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

Backus, D. K., and P. J. Kehoe. "International Evidence on the Historical Properties of Business Cycles." *American Economic Review*, 82 (1992), 864–888.

Balvers, R. J., and D. W. Mitchell. "Autocorrelated Returns and Optimal Intertemporal Portfolio Choice." *Management Science*, 43 (1997), 1537–1551.

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

Baxter, M., and R. King. "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series." *Review of Economics and Statistics*, 81 (1999), 575–593.

Brandt, M. "Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach." *Journal of Finance*, 54 (1999), 1609–1645.

Brandt, M. W., and P. Santa-Clara. "Dynamic Portfolio Selection by Augmenting the Asset Space." *Journal of Finance*, 61 (2006), 2187–2217.

Brennan, M. J.; E. S. Schwartz; and R. Lagnado. "Strategic Asset Allocation." *Journal of Economic Dynamics and Control*, 21 (1997), 1377–1403.

Campbell, J. Y.; Y. L. Chan; and L. M. Viceira. "A Multivariate Model of Strategic Asset Allocation." *Journal of Financial Economics*, 67 (2003), 41–80.

Campbell, J. Y, and L. Viceira. "Consumption and Portfolio Decisions When Expected Returns Are Time Varying." *Quarterly Journal of Economics*, 114 (1999), 433–495.

Campbell, J. Y., and L. M. Viceira. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford, UK: Oxford University Press (2003).

Christiano, L., and T. Fitzgerald. "The Band Pass Filter." *International Economic Review*, 44 (2003), 435–465.

Cochrane, J. "New Facts in Finance." *Economic Perspectives*, 23 (1999), 36–38.

Cochrane, J. "The Dog That Did Not Bark: A Defense of Return Predictability." *Review of Financial Studies*, 21 (2008), 1533–1577.

Cremers, K. "Stock Return Predictability: A Bayesian Model Selection Perspective." *Review of Financial Studies*, 15 (2002), 1223–1249.

Doornik, J. *An Object-Oriented Matrix Language Ox 6*. London, UK: Timberlake Consultants Press (2009).

Doornik, J., and M. Ooms. *Introduction to Ox: An Object-Oriented Matrix Language*. London, UK: Timberlake Consultants Press (2007).

Fama, E., and K. French. "Dividend Yields and Expected Stock Returns." *Journal of Financial Economics*, 22 (1988), 3–25.

Fama, E., and K. French. "Business Conditions and Expected Returns on Stocks and Bonds." *Journal of Financial Economics*, 25 (1989), 23–49.

Ferson, W. E., and C. R. Harvey. "Sources of Predictability in Portfolio Returns." *Financial Analysts Journal*, 47 (1991), 49–56.

Ferson, W. E., and C. R. Harvey. "Explaining the Predictability in Asset Returns." *Research in Finance*, 11 (1993), 65–106.

Ferson, W. E., and C. R. Harvey. "Fundamental Determinants of National Equity Market Returns: A Perspective on Conditional Asset Pricing." *Journal of Banking and Finance*, 21 (1998), 1625–1665.

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

Handa, P. "Does Stock Return Predictability Imply Improved Asset Allocation and Performance? Evidence from the US Stock Market (1954–2002)." *Journal of Business*, 79 (2006), 2423–2468.

Hansen, L. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica: Journal of the Econometric Society*, 50 (1982), 1029–1054.

Harvey, A. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge, UK: Cambridge University Press (1990).

Harvey, A., and A. Jaeger. "Detrending, Stylized Facts and the Business Cycle." *Journal of Applied Econometrics*, 8 (1993), 231–247.

Hodrick, R., and E. Prescott. "Postwar U.S. Business Cycles: An Empirical Investigation." *Journal of Money, Credit and Banking*, 29 (1997), 1–16.

Jegadeesh, N., and S. Titman. “Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency.” *Journal of Finance*, 48 (1993), 65–91.

Jobson, J., and B. Korkie. “Performance Hypothesis Testing with the Sharpe and Treynor Measures.” *Journal of Finance*, 36 (1981), 889–908.

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

Keim, D. B., and R. F. Stambaugh. “Predicting Returns in the Stock and Bond Markets.” *Journal of Financial Economics*, 17 (1986), 357–390.

Koijen, R. S.; J. C. Rodriguez; and A. Sbuelz. “Momentum and Mean Reversion in Strategic Asset Allocation.” *Management Science*, 55 (2009), 1199–1213.

Leser, C. “A Simple Method of Trend Construction.” *Journal of the Royal Statistical Society. Series B (Methodological)*, 23 (1961), 91–107.

Lettau, M., and S. Ludvigson. “Expected Returns and Expected Dividend Growth.” *Journal of Financial Economics*, 76 (2005), 583–626.

Lettau, M., and S. Van Nieuwerburgh. “Reconciling the Return Predictability Evidence.” *Review of Financial Studies*, 21 (2008), 1607–1652.

Lewellen, J. “Momentum and Autocorrelation in Stock Returns.” *Review of Financial Studies*, 15 (2002), 533–564.

Lynch, A. W. “Portfolio Choice and Equity Characteristics: Characterizing the Hedging Demands Induced by Return Predictability.” *Journal of Financial Economics*, 62 (2001), 67–130.

Memmel, C. “Performance Hypothesis Testing with the Sharpe Ratio.” *Finance Letter*, 1 (2003), 21–23.

Pettenuzzo, D., and A. Timmermann. “Predictability of Stock Returns and Asset Allocation under Structural Breaks.” *Journal of Econometrics*, 164 (2011), 60–78.

Shanken, J., and A. Tamayo. “Payout Yield, Risk, and Mispricing: A Bayesian Analysis.” *Journal of Financial Economics*, 105 (2012), 131–152.

Tamayo, A. “Stock Return Predictability, Conditional Asset Pricing Models and Portfolio Selection.” Working Paper, London Business School (2002).

Van Binsbergen, J., and R. Koijen. “Predictive Regressions: A Present-Value Approach.” *Journal of Finance*, 65 (2010), 1439–1471.

Wachter, J., and M. Warusawitharana. “Predictable Returns and Asset Allocation: Should a Skeptical Investor Time the Market?” *Journal of Econometrics*, 148 (2009), 162–178.

Wachter, J. A. “Portfolio and Consumption Decisions under Mean-Reverting Returns: An Exact Solution for Complete Markets.” *Journal of Financial and Quantitative Analysis*, 37 (2002), 63–91.