EGRIE Keynote Address

The Changing Role of Nominal Government Bonds in Asset Allocation*

John Y. Campbell

Department of Economics, Littauer Center, Harvard University, Cambridge MA 02138, U.S.A. and NBER.

The covariance between nominal bonds and stocks has varied considerably over recent decades and has even switched sign. It has been predominantly positive in periods such as the late 1970s and early 1980s when the economy has experienced supply shocks and the central bank has lacked credibility. It has been predominantly negative in periods such as the 2000s when investors have feared weak aggregate demand and deflation. Nominal bonds are attractive to short-term equity investors when these bonds are negatively correlated with stocks, as has been the case during the 2000s and especially during the downturn of 2007–2008. They are attractive to conservative long-term investors when long-term inflationary expectations are stable, for then these bonds are close substitutes for inflation-indexed bonds that are riskless in the long term.

The Geneva Risk and Insurance Review (2009) **34**, 89–104. doi:10.1057/grir.2009.7

Keywords: Mean-variance analysis; long-term investing; time-varying risk

Introduction

How should households saving for retirement allocate their portfolios across different asset classes such as stocks, nominal government bonds, inflation-indexed government bonds, and money market instruments or “cash”? Conventional analysis of this question assumes that broad asset classes have stable risks, which can be measured by looking at the covariances of asset classes over long periods of history. Even research that emphasizes the distinction between the risks faced by short-term investors and those faced by long-term investors (Campbell and Viceira’s (2005) “term structure of the risk-return tradeoff”) tends to assume that this term structure is constant over time.

*This paper is based on the Geneva Risk Economics Lecture delivered in Toulouse, France to the European Group of Risk and Insurance Economists on 16 September 2008.

In recent years it has become clear that the relative risks of nominal government bonds and stocks are not constant over time. I will illustrate the point using U.S. data, but similar patterns are evident in other countries as well. Figure 1, taken from Viceira (2007), shows one measure of the risk of bonds relative to stocks, the beta or regression coefficient of daily nominal 10-year zero-coupon Treasury bond returns on stock returns, measured within a rolling 3-month window from July 1962 to December 2003. The figure shows high-frequency variation from one quarter to the next in the realized beta of bonds on stocks, much of which is unpredictable noise. It also shows low-frequency movements in the beta, which was close to zero but slightly positive on average in the 1960s and early 1970s, was considerably higher with an average of about 0.2 in the 1980s and again in the mid-1990s, and turned negative in the late 1990s.

The negative average beta of nominal Treasury bonds has persisted throughout the current decade. Figure 2 plots the same beta coefficient over the period from June 2002 through April 2008. The average is clearly negative, and particularly so in the downturns of the early 2000s and 2007–2008. Campbell *et al.* (2009b) and Donovan *et al.* (2008) report similar results using recent data from both the U.S. and the U.K. The latter paper uses both asymptotic theory of Barndorff-Nielsen and Shephard (2004) and bootstrap simulations to show that the sign switches in realized betas are statistically significant.

The beta of nominal bonds with stocks measures the risk that a small bond investment, financed by short-term borrowing, adds to a portfolio initially

Figure 1. CAPM beta of bonds (July 1962–December 2003). Realized beta of bonds based on 3-months of daily returns on stocks and bonds.

Source: Viceira (2007).

Figure 2. CAPM beta of bonds (June 2002–April 2008).

invested in equities. When this beta is positive, bonds are incrementally risky and will only be attractive to equity investors if they offer a positive term premium (that is, a positive expected excess return over cash). When the beta is negative, however, bonds act as a hedge against equity risk and may be held for this reason even if the term premium is zero or negative. Thus time-variation in the beta of bonds with stocks can have profound implications for asset allocation.

Both academics and investment practitioners have changed their attitudes towards nominal bonds over the decades, mirroring the low-frequency movements in bond risks illustrated in Figures 1 and 2. In the late 1970s and early 1980s, the Wall Street economist Henry Kaufman rose to prominence by emphasizing the risk that inflation posed to bond investors, while academic research emphasized that bonds should offer a large-term premium to compensate for this inflation risk exposure. This view influenced the decision of the U.K. government to issue inflation-indexed bonds in the early 1980s, followed much later by the U.S. government in 1997. By the 2000s, in contrast, nominal bonds were seen as relatively safe investments, and even hedges against slow growth accompanied by deflation of the sort that Japan experienced in the 1990s.

In this paper I argue that investors need to understand and respond to variation over time in the relative risks of nominal government bonds and stocks. I begin in the following section by surveying recent work that models this variation. I indicate fruitful directions for future research on this topic. In the subsequent section, I explore implications for optimal asset allocation. The last section concludes.

Modeling time-varying bond risk

The importance of inflation

Inflation is relevant for investors in nominal government bonds because these investors are promised fixed nominal payments, not fixed real payments. The greater is realized inflation over the life of the bond, the lower the real return on the investment. Therefore nominal bond prices fall when expected inflation increases, and movements in expected inflation are a major source of short-term volatility in bond returns.

Figure 3 shows that there have been changes in the covariance between *realized* inflation and stock returns, mirroring the changes in the covariance between nominal bond and stock returns shown in Figure 1. The figure works with deflation, the negative of inflation, because deflation is positively related to nominal bond returns; and because consumer prices are only measured at a

Figure 3. CAPM beta of deflation (3-year rolling window of monthly shocks to log inflation and stock returns).

monthly frequency, it uses a 3-year window of monthly data rather than a 3-month window of daily data to calculate the realized beta of deflation with stock returns. The same low-frequency variations that were visible in Figure 1 appear in Figure 3 as well.

One can also look at the covariance between *expected* inflation and stock returns. Campbell *et al.* (2009b) measure breakeven inflation, the difference in yield between nominal and inflation-indexed Treasury bonds of the same maturity. In normal market conditions, breakeven inflation is a reasonable measure of expected inflation, although technical dislocation in the bond market in the fall of 2008 created unusual variations in breakeven inflation that may not accurately indicate market participants' expectations of inflation. Campbell, Shiller, and Viceira show that daily movements in breakeven inflation have been positively correlated with stock returns during the 2000s, especially in the early part of the decade and the 2007–2008 downturn. Thus breakeven deflation has been negatively correlated with stock returns during this period, helping to explain the negative beta of nominal bonds with stocks.

Macroeconomic models can be used to understand why the covariance of inflation with the stock market might change over time. Stock prices are procyclical, so inflation is likely to covary positively with stock prices if it is procyclical, covarying positively with the real economy. Traditional Keynesian models with a stable Phillips Curve imply that inflation is procyclical, as strong aggregate demand drives up product prices. If the Phillips Curve shifts outward, however, as famously occurred in the 1970s, then inflation increases even though the economy is weak. Such “stagflation” can occur if the economy is subjected to supply shocks or if monetary policy loses credibility with the public, allowing long run expected inflation to increase. New Keynesian models use an expectations-augmented Phillips Curve to capture this effect.

The lesson of this analysis is that periods with supply shocks or poor central bank credibility, such as the 1970s and early 1980s, are likely to have countercyclical inflation (procyclical deflation) and a positive beta of nominal bonds with stocks; while periods with demand shocks and credible monetary policy, such as the 1950s and 2000s, are more likely to have procyclical inflation (countercyclical deflation) and a negative beta of nominal bonds with stocks.

A formal model

The evidence I have presented implies that a satisfactory model of nominal bond pricing must have three properties. First, it must allow for changes over time in the risks of nominal bonds. Second, it must allow the covariance between bond and stock returns to switch sign. Third, the changing risks of nominal bonds should be linked to the behaviour of inflation.

It is not straightforward to build a model with all three of these properties. Many simple models of changing bond risk premia are driven by a single time-varying volatility process, either for the real interest rate (Cox *et al.*, 1985) or for the stochastic discount factor (SDF). Models of this sort scale covariances up and down but do not allow them to switch sign. More generally, it is difficult to allow for sign switches in covariances while remaining within the tractable affine class of models in which log bond yields are linear in state variables (Dai and Singleton, 2002; Duffee, 2002). Also, many bond pricing models are not fully explicit about the distinction between real and nominal quantities.

Campbell, Sunderam, and Viceira (CSV, 2009c) write down a simple model that does meet these three criteria. Their model is a traditional affine model of the real yield curve, augmented with a time-varying covariance between inflation and the real economy. The resulting nominal term structure model is linear-quadratic in macroeconomic state variables.¹

The real economy, real interest rates, and the stock market

CSV begin by assuming that the log of the real SDF $m_{t+1} = \log(M_{t+1})$ follows a linear-quadratic, conditionally heteroskedastic process:

$$-m_{t+1} = x_t + \frac{\sigma_m^2}{2} z_t^2 + z_t \varepsilon_{m,t+1}, \quad (1)$$

where both x_t and z_t follow standard AR(1) processes. Given homoskedasticity of underlying shocks ε , the log real SDF is conditionally heteroskedastic, with

$$\text{Var}_t(m_{t+1}) = z_t^2.$$

The state variable z_t drives the time-varying volatility of the SDF or, equivalently, the price of aggregate market risk or maximum Sharpe ratio in the economy. It can be understood as a measure either of changing risk aversion (Campbell and Cochrane, 1999, Bekaert *et al.*, 2005), or of changing volatility in the real economy (Bansal and Yaron, 2004).

It is straightforward to show that the one-period real interest rate equals the state variable x_t , and the whole term structure of real interest rates is linear in the two real state variables x_t and z_t . To bring stock returns into the model,

¹ Other linear-quadratic term structure models include Beaglehole and Tenney (1991), Constantinides (1992), and Ahn *et al.* (2002). Duffie and Kan (1996) point out that linear-quadratic models can often be rewritten as affine models if we allow the state variables to be bond yields rather than macroeconomic fundamentals. Buraschi *et al.* (2008) also expand the state space to obtain an affine model in which correlations can switch sign.

CSV write down a reduced form equation expressing shocks to realized stock returns as a linear combination of shocks to the real interest rate and shocks to the log SDF. This implies that the equity premium, like all other risk premia in the model, is proportional to risk aversion z_t . It depends not only on the direct sensitivity of stock returns to the SDF, but also on the sensitivity of stock returns to the real interest rate and the covariance of the real interest rate with the SDF.

Inflation and nominal interest rates

To price nominal bonds, CSV specify a model for inflation. They assume that log inflation π_t follows a linear-quadratic conditionally heteroskedastic process:

$$\pi_{t+1} = \lambda_t + \xi_t + \frac{\sigma_\pi^2}{2} \psi_t^2 + \psi_t \varepsilon_{\pi,t+1}, \quad (2)$$

where expected log inflation is the sum of two components, a permanent component λ_t and a transitory component ξ_t , both driven by underlying shocks that are also scaled by the state variable ψ_t .

The inclusion of two components of expected inflation gives the model the flexibility it needs to fit simultaneously persistent shocks to both real interest rates and expected inflation. This flexibility is necessary because both realized inflation and the yields of long-dated inflation-indexed bonds move persistently, which suggests that both expected inflation and the real interest rate follow highly persistent processes. At the same time, short-term nominal interest rates exhibit more variability than long-term nominal interest rates, which suggests that a rapidly mean-reverting state variable must also drive the dynamics of nominal interest rates.

The state variable, ψ_t , which multiplies the underlying shocks that drive realized and expected inflation, is assumed to follow a homoskedastic AR(1) process with a non-zero mean. This specification implies that the conditional volatility of inflation is time varying, as in the original ARCH model of Engle (1982). The novel feature of the specification is that ψ_t can change sign. The sign of ψ_t does not affect the variances of expected or realized inflation or the covariance between them, because these moments depend on the square ψ_t^2 . However, the sign of ψ_t does determine the sign of the covariance between expected and realized inflation, on the one hand, and real economic variables, on the other hand. Thus it can track the changes in covariances illustrated in Figures 1 and 2.

CSV show that under these assumptions the log nominal short rate is a linear-quadratic function of the state variables, and this property carries over to the entire zero-coupon nominal term structure. The log price of a n -period

zero-coupon nominal bond can be written as a linear function of the state variables x_t , z_t , λ_t , ξ_t , and ψ_t , and the squares and cross-product z_t^2 , ψ_t^2 , and $z_t\psi_t$.

CSV estimate the model using a nonlinear “unscented” Kalman filter (Wan and van der Merwe, 2001) to construct the likelihood function. They find that the term structure is driven by shocks to the permanent component of expected inflation λ_t , which move the entire yield curve up and down (“level” shocks in the terminology of fixed-income practitioners), shocks to real interest rates x_t and the temporary component of expected inflation ξ_t , which move short rates more than long rates (“slope” shocks), and shocks to risk aversion z_t and the covariance of real and nominal magnitudes ψ_t , which alter risk premia and the concavity of the yield curve (“curvature” shocks). The last two shocks drive risk premia on nominal bonds, which are approximately proportional to the product $z_t\psi_t$. In this way the model helps to explain the empirical association between concavity of the yield curve and excess bond returns, noted by Cochrane and Piazzesi (2005) among others.

Extending the model

The work of CSV can be extended in several directions. One limitation is that the affine structure of the real side of the model implies a constant covariance between inflation-indexed bonds and stocks. Campbell *et al.* (2009b) show that both Treasury inflation-protected securities (TIPS) in the U.S. and inflation-indexed gilts in the U.K. have moved more negatively with stocks during the downturns of the early 2000s and 2007–2008 than they did in the mid-2000s or (in the U.K.) the 1990s. To capture this they introduce a state variable that moves the covariance of real interest rates with the SDF, a real-side analog to the nominal variable ψ_t .

Ultimately, one would like to have a deeper structural understanding of the origin of these fluctuations in covariances. It should be possible to achieve this by writing down a New Keynesian macroeconomic model and allowing some of the parameters, including perhaps the volatilities of shocks and the parameters describing monetary policy, to vary over time as in Clarida *et al.* (2000). This raises the exciting possibility that one can use the changing covariances between stocks and real and nominal government bonds to learn about the nature of the underlying macroeconomic regime.

Asset allocation with time-varying bond risk

How should investors respond to changes over time in the covariance between nominal bonds and stocks? It is important at the outset to distinguish between

short-term investors, who are concerned with the distribution of invested wealth a quarter or a year ahead, and long-term investors, who measure risk by the distribution of wealth many years ahead or even by the sustainable consumption stream that wealth can support.

The changing role of bonds for short-term investors

Short-term investors have an almost entirely safe asset available in the form of Treasury bills, whose nominal return is guaranteed and whose real return has minimal variability given that inflation is highly predictable over a quarter and even over a year. It follows that short-term investors hold long-term bonds not for safety, but either for their expected excess return (the “speculative motive”) or for their ability to hedge the risks of other assets such as equities (the “hedging motive”).

The standard mean-variance analysis of Markowitz (1952) can be used to evaluate the role of bonds in risky portfolios for short-term investors. Short-term mean-variance investors invest in a unique tangency portfolio of risky assets, combining this with Treasury bills in proportions that depend on the risk aversion of each investor. If the two available risky assets are nominal Treasury bonds and the aggregate U.S. stock market, the weight of bonds in the tangency portfolio depends on the mean excess returns of bonds and stocks, their variances, and the covariance between them.

If in addition mean excess returns and the variance of stock returns are reasonably stable over time, then the role of nominal bonds in the tangency portfolio depends primarily on their volatility and their covariance with the stock market. When bonds are positively correlated with stocks, they have a relatively small weight in the tangency portfolio and that portfolio is quite volatile. When bonds are negatively correlated with stocks, they have a larger weight in the tangency portfolio because of their ability to hedge stock market risk. The tangency portfolio is also more stable and has a higher Sharpe ratio (return per unit of risk).

These properties are illustrated in Figure 4, which shows the ratio of stocks to bonds in the tangency portfolio implied by Campbell *et al.*’s (2009c) filtered estimates of their term structure model. Although expected returns do vary in the CSV model, they do not move enough to offset the effects of changing risks on the composition of the tangency portfolio. In the early 1980s the tangency portfolio is dominated by stocks and is correspondingly volatile, whereas in the 1950s, 1960s, and 2000s, bonds play a dominant role with a stock-bond ratio less than one. At such times the stability of the tangency portfolio encourages aggressive investors to use leverage. This suggests that the negative correlation between nominal bonds and stocks in the 2000s may have contributed to the increased use of leverage during the credit boom of the mid-2000s.

Figure 4. Stock/bond ratio in the tangency portfolio of the Campbell *et al.* (2009c) model.

The changing role of bonds for long-term investors

Campbell and Viceira (2001, 2002) have emphasized that long-term bonds play a more important role for long-term investors. For these investors, Treasury bills are not safe because they must be rolled over at uncertain future interest rates. An investor who seeks safety at a fixed long horizon can achieve it by buying a zero-coupon inflation-indexed bond of the given maturity, and an investor who seeks a safe consumption stream that is indefinitely sustainable can achieve it by buying an inflation-indexed perpetuity. If inflation-indexed bonds are not available, long-term investors must combine other assets, including Treasury bills, nominal bonds, and stocks, to minimize their risk.

Campbell and Viceira (2005) specifically show how to calculate a global minimum-variance (GMV) portfolio at any investment horizon, using a vector autoregressive (VAR) model to capture changes over time in real interest rates and expected bond and stock returns. Their analysis assumes that the covariance matrix of shocks to the VAR is constant over time; thus they do

not consider the phenomenon of changing covariances discussed in this paper. Figure 5, taken from their paper, shows how the GMV portfolio weights of Treasury bills, 5-year nominal Treasury bonds, and stocks change with the investment horizon. The figure is based on a covariance matrix of shocks that is estimated over Campbell and Viceira's full sample period 1953–2002. At short horizons, the GMV portfolio is dominated by Treasury bills, with modest short positions in stocks and bonds to hedge against inflation shocks that lower real bill returns and also lower the prices of stocks and bonds. At longer horizons, the rollover risk of Treasury bills becomes more important, so nominal Treasury bonds become the dominant asset in the GMV portfolio.

Figures 6, 7, and 8 show how these conclusions are altered by estimating the VAR covariance matrix over three different 5-year periods, chosen to illustrate three different regimes in asset markets. In the mid-1950s (1953–1957) real interest rates were extremely stable, so there was little rollover risk in Treasury bills, which remain the dominant asset in the GMV portfolio out to a 100-year investment horizon (Figure 6).

Figure 5. Composition of Global Minimum Variance portfolio (1953.Q2–2002.Q4).

In the mid-1980s (1983–1987), real interest rates were volatile implying that Treasury bills were not safe long-term assets. At the same time there was great uncertainty about inflationary conditions, so nominal Treasury bonds were not similar to inflation-indexed bonds and did not offer safe long-term returns. In this period, equities play a major role in the long-term GMV portfolio and short positions in nominal bonds, which were positively correlated with stocks at this time, are used to hedge equity risk (Figure 7).

Finally, around the turn of the millennium (1998–2002), real interest rates were volatile but long-term expectations of inflation were stable. This implies that nominal Treasury bonds are extremely similar to inflation-indexed bonds and play a dominant role in the long-term GMV portfolio (Figure 8). Campbell *et al.* (2009b) show that TIPS have had a correlation with nominal Treasuries close to one for much of this decade, supporting the plausibility of this finding.

One caveat about the long-term GMV analysis should be mentioned here. Campbell and Viceira's (2005) methodology assumes that a portfolio must be chosen once and for all at the start of the investment horizon, without allowing

Figure 6. Composition of Global Minimum Variance portfolio (1953–1957).

Figure 7. Composition of Global Minimum Variance portfolio (1983–1987).

rebalancing to respond to changing investment opportunities. However, a full intertemporal analysis along the lines of Merton (1973) delivers similar results in the empirical implementation of Campbell *et al.* (2003).

Conclusion

Traditional asset allocation analysis assumes that asset classes have stable risks that can be estimated from long-term historical data. Even sophisticated approaches that recognize changes over time in expected returns, and the resulting differences in the risks perceived by short-term and long-term investors typically ignore the fact that risks may also change over time.

When nominal bonds are included in an asset allocation exercise, as is almost always the case, the assumption of constant risks is dangerously misleading. The covariance between nominal bonds and stocks has varied considerably over recent decades and has even switched sign. It has been predominantly

Figure 8. Composition of Global Minimum Variance portfolio (1998–2002).

positive in periods such as the late 1970s and early 1980s when the economy has experienced supply shocks and the central bank has lacked credibility. It has been predominantly negative in periods such as the 2000s when investors have feared weak aggregate demand and deflation.

Nominal bonds are attractive to short-term equity investors when these bonds are negatively correlated with stocks, as has been the case during the 2000s and especially during the downturn of 2007–2008. They are attractive to conservative long-term investors when long-term inflationary expectations are stable, for then these bonds are close substitutes for inflation-indexed bonds that are riskless in the long term. At present, nominal bonds therefore play an important role in asset allocation even if they offer a small or negative term premium over Treasury bills.

The demand for nominal bonds in asset allocation can, however, change rapidly if the regime changes. If investors come to fear stagflation, bonds' ability to hedge against deflation will no longer be so attractive, and the correlation between bonds and stocks may switch sign once again. If

inflationary expectations destabilize, nominal bonds are no longer close substitutes for inflation-indexed bonds and are less appealing for conservative long-term portfolios. Both investors and fiscal and monetary authorities should pay close attention to changing covariances among nominal bonds, inflation-indexed bonds, and stocks as a guide to asset allocation and an indicator of the state of the economy.

The importance of the macroeconomic regime for asset allocation applies beyond the specific example discussed in this paper. Many other asset classes, including foreign currencies (Campbell *et al.*, 2009a), real estate, and commodities, also have risks that are likely to vary with the economic environment. Much as investors might wish to choose portfolios based on mechanical processing of historical data, asset allocation cannot be conducted without forming a view about the structure of the economy and the relative magnitudes of the shocks that impinge upon it.

Acknowledgements

I acknowledge the support of the U.S. Social Security Administration (SSA) through grant #10-M-98363-1-01 to the National Bureau of Economic Research (NBER) as part of the SSA Retirement Research Consortium. The findings and conclusions expressed are solely those of the author and do not represent the views of SSA, any agency of the federal government, or the NBER. I am grateful to an anonymous referee for comments, to Carolin Pflueger for able research assistance, and to Robert Shiller, Adi Sunderam, and especially Luis Viceira for their collaboration on research underlying this paper.

References

- Ahn, D.-H., Dittmar, R.F. and Gallant, A.R. (2002) 'Quadratic term structure models: Theory and evidence', *Review of Financial Studies* **15**: 243–288.
- Bansal, R. and Yaron, A. (2004) 'Risks for the long run: A potential resolution of asset pricing puzzles', *Journal of Finance* **59**: 1481–1509.
- Barndorff-Nielsen, O. and Shephard, N. (2004) 'Econometric analysis of realized covariation: High frequency based covariance, regression, and correlation in financial econometrics', *Econometrica* **72**: 885–925.
- Beaglehole, D.R. and Tenney, M.S. (1991) 'General solutions of some interest rate-contingent claim pricing equations', *Journal of Fixed Income* **1**(September): 69–83.
- Bekaert, G., Engstrom, E. and Grenadier, S. (2005) *Stock and bond returns with moody investors*, unpublished paper, Columbia University, University of Michigan, and Stanford University.
- Buraschi, A., Cieslak, A. and Trojani, F. (2008) *Correlation risk and the term structure of interest rates*, unpublished paper, Imperial College London and University of St. Gallen.
- Campbell, J.Y. and Cochrane, J.H. (1999) 'By force of habit: A consumption-based explanation of aggregate stock market behavior', *Journal of Political Economy* **107**: 205–251.
- Campbell, J.Y., Lewis Chan, Y. and Viceira, L.M. (2003) 'A multivariate model of strategic asset allocation', *Journal of Financial Economics* **67**: 41–80.

- Campbell, J.Y., Serfaty-de Medeiros, K. and Viceira, L.M. (2009a) *Global currency hedging*, NBER working paper No. 13088, forthcoming *Journal of Finance*.
- Campbell, J.Y., Shiller, R.J. and Viceira, L.M. (2009b) "Understanding inflation-indexed bond markets", *Brookings Papers on Economic Activity* (Spring): 79–120.
- Campbell, J.Y., Sunderam, A. and Viceira, L.M. (2009c) *Inflation bets or deflation hedges? The changing risks of nominal bonds*, NBER working paper No. 14701.
- Campbell, J.Y. and Viceira, L.M. (2001) 'Who should buy long-term bonds?' *American Economic Review* **91**: 99–127.
- Campbell, J.Y. and Viceira, L.M. (2002) *Strategic Asset allocation: Portfolio Choice for Long-Term Investors*, New York: Oxford University Press.
- Campbell, J.Y. and Viceira, L.M. (2005) 'The term structure of the risk-return tradeoff', *Financial Analysts Journal* **61**(January/February): 34–44.
- Clarida, R., Gali, J. and Gertler, M. (2000) 'Monetary policy rules and macroeconomic stability: Evidence and some theory', *Quarterly Journal of Economics* **115**: 147–180.
- Cochrane, J.H. and Piazzesi, M. (2005) 'Bond risk premia', *American Economic Review* **95**: 138–160.
- Constantinides, G.M. (1992) 'A theory of the nominal term structure of interest rates', *Review of Financial Studies* **5**: 531–552.
- Cox, J.C., Ingersoll, J.E. and Ross, S.A. (1985) 'An intertemporal general equilibrium model of asset prices', *Econometrica* **53**: 363–384.
- Dai, Q. and Singleton, K. (2002) 'Expectations puzzles, time-varying risk premia, and affine models of the term structure', *Journal of Financial Economics* **63**: 415–442.
- Donovon, P., Gonçalves, S. and Meddahi, N. (2008) *Bootstrapping realized multivariate volatility measures*, unpublished paper, Université de Montréal.
- Duffee, G. (2002) 'Term premia and interest rate forecasts in affine models', *Journal of Finance* **57**: 405–443.
- Duffie, D. and Kan, R. (1996) 'A yield-factor model of interest rates', *Mathematical Finance* **6**: 379–406.
- Engle, R.F. (1982) 'Autoregressive conditional heteroskedasticity with estimates of the variance of UK Inflation', *Econometrica* **50**: 987–1008.
- Markowitz, H. (1952) 'Portfolio selection', *Journal of Finance* **7**: 77–91.
- Merton, R.C. (1973) 'An intertemporal capital asset pricing model', *Econometrica* **41**: 867–887.
- Viceira, L.M. (2007) *Bond risk, bond return volatility, and the term structure of interest rates*, unpublished paper, Harvard Business School.
- Wan, E.A. and van der Merwe, R. (2001) 'The unscented Kalman filter', Chapter 7 in S. Haykin (ed) *Kalman Filtering and Neural Networks*, New York: Wiley.

About the Author

John Y. Campbell is the Morton L. and Carole S. Olshan Professor of Economics and Chair of the Department of Economics at Harvard University. He is also a research associate of the National Bureau of Economic Research, a board member of Harvard Management Company, and a partner at Arrowstreet Capital, LP.