

Understanding Inflation-Indexed Bond Markets

John Y. Campbell, Robert J. Shiller, and Luis M. Viceira¹

First draft: February 2009

This version: May 2009

¹Campbell: Department of Economics, Littauer Center, Harvard University, Cambridge MA 02138, and NBER. Email john_campbell@harvard.edu. Shiller: Cowles Foundation, Box 208281, New Haven CT 06511, and NBER. Email robert.shiller@yale.edu. Viceira: Harvard Business School, Boston MA 02163 and NBER. Email lviceira@hbs.edu. Campbell and Viceira's research was supported by the U.S. Social Security Administration through grant #10-M-98363-1-01 to the National Bureau of Economic Research as part of the SSA Retirement Research Consortium. The findings and conclusions expressed are solely those of the authors and do not represent the views of SSA, any agency of the Federal Government, or the NBER. We are grateful to Carolin Pflueger for exceptionally able research assistance, to Mihir Worah and Gang Hu of PIMCO, Derek Kaufman of Citadel, and Albert Brondolo, Michael Pond, and Ralph Segreti of Barclays Capital for their help in understanding TIPS and inflation derivatives markets and the unusual market conditions in the fall of 2008, and to Barclays Capital for providing data. An earlier version of the paper was presented at the Brookings Panel on Economic Activity, April 2-3, 2009. We acknowledge the helpful comments of panel members and our discussants, Rick Mishkin and Jonathan Wright.

Abstract

This paper explores the history of inflation-indexed bond markets in the US and the UK. It documents a massive decline in long-term real interest rates from the 1990's until 2008, followed by a sudden spike in these rates during the financial crisis of 2008. Breakeven inflation rates, calculated from inflation-indexed and nominal government bond yields, stabilized until the fall of 2008, when they showed dramatic declines. The paper asks to what extent short-term real interest rates, bond risks, and liquidity explain the trends before 2008 and the unusual developments in the fall of 2008. Low inflation-indexed yields and high short-term volatility of inflation-indexed bond returns do not invalidate the basic case for these bonds, that they provide a safe asset for long-term investors. Governments should expect inflation-indexed bonds to be a relatively cheap form of debt financing going forward, even though they have offered high returns over the past decade.

1 Introduction

In recent years government inflation-indexed bonds have become available in a number of countries and have provided a fundamentally new instrument for use in retirement saving. Because expected inflation varies over time, long-term nominal Treasury bonds are not safe in real terms; and because short-term real interest rates vary over time, Treasury bills are not safe assets for long-term investors. Inflation-indexed bonds fill this gap by offering a truly riskless long-term investment (Campbell and Shiller 1996, Campbell and Viceira 2001, 2002, Brennan and Xia 2002, Campbell, Chan, and Viceira 2003, Wachter 2003).

The UK government issued inflation-indexed bonds in the early 1980's, and the US government followed suit by issuing Treasury inflation-protected securities (TIPS) in 1997. Inflation-indexed government bonds are also available in many other countries including Canada, France, and Japan. These bonds are now widely accepted financial instruments. However, their history raises some new puzzles that deserve investigation.

First, given that the real interest rate is determined by the marginal product of capital in the long run, one might expect inflation-indexed yields to be extremely stable over time. But during the 1990's, 10-year inflation-indexed yields averaged about 3.5% in the UK (Barr and Campbell 1997), and exceeded 4% in the US around the turn of the millennium, whereas in the mid-2000's they both averaged below 2% and bottomed out at around 1% in early 2008 before spiking up above 3% in late 2008. The massive decline in long-term real interest rates from the 1990's to the 2000's is one puzzle, and the instability in 2008 is another.

Second, in recent years inflation-indexed bond prices have tended to move opposite stock prices, so that these bonds have a negative "beta" with the stock market and can be used to hedge equity risk. This has been even more true of nominal bond prices, although nominal bonds behaved very differently in the 1970's and 1980's (Campbell, Sunderam, and Viceira 2009). The origin of the negative beta for inflation-indexed bonds is not well understood.

Third, given integrated world capital markets, one might expect that inflation-indexed bond yields would be similar around the world. But this is not always the case. Around the year 2000, the yield gap between US and UK inflation-indexed bonds was over 2 percentage points, although it has since converged. In January 2008,

while 10-year yields were similar in the US and the UK, there were still important differentials across countries, with yields ranging from 1.1% in Japan to almost 2.0% in France. Yield differentials were even larger at long maturities, with UK yields well below 1% and French yields well above 2%.

To understand these phenomena, it is useful to distinguish three major influences on inflation-indexed bond yields: current and expected future short-term real interest rates; differences in expected returns on long-term and short-term real bonds caused by risk premia (which can be negative if inflation-indexed bonds are valuable hedges); and differences in expected returns on long-term and short-term bonds caused by liquidity premia or technical factors that segment the bond markets. The expectations hypothesis of the term structure, applied to real interest rates, states that only the first influence is time-varying while the other two are constant. However there is considerable evidence against this hypothesis for nominal Treasury bonds, so it is important to allow for the possibility that risk and liquidity premia are time-varying.

Undoubtedly the path of real interest rates is a major influence on inflation-indexed bond yields. Indeed, before TIPS were issued Campbell and Shiller (1996) argued that one could anticipate how their yields would behave by applying the expectations hypothesis of the term structure to real interest rates. A first goal of this paper is to compare the history of inflation-indexed bond yields with the implications of the expectations hypothesis, and to understand how shocks to short-term real interest rates are transmitted along the real yield curve.

Risk premia on inflation-indexed bonds can be analyzed by applying theoretical models of risk and return. Two leading paradigms deliver useful insights. The consumption-based paradigm implies that risk premia on inflation-indexed bonds over short-term debt are negative if these bonds covary negatively with consumption, which will be the case if consumption growth rates are persistent (Backus and Zin 1994, Campbell 1986, Gollier 2005, Piazzesi and Schneider 2006, Wachter 2006), while the CAPM paradigm implies that inflation-indexed risk premia are negative if inflation-indexed bond prices covary negatively with stock prices. The second paradigm has the advantage that it is easy to track the covariance of inflation-indexed bonds and stocks using high-frequency data on their prices, in the manner of Viceira (2007) and Campbell, Sunderam, and Viceira (2009).

Finally, it is important to take seriously the effects of institutional factors on inflation-indexed bond yields. Plausibly, the high TIPS yields in the first few years after their introduction were caused by slow development of mutual funds and other

indirect investment vehicles. Currently, long-term inflation-indexed yields in the UK may be depressed by strong demand from UK pension funds. The volatility of TIPS yields in the fall of 2008 appears to have resulted in part from the unwinding of large institutional positions after the failure of Lehman Brothers. These institutional influences on yields can alternatively be described as liquidity, market segmentation, or demand and supply effects (Greenwood and Vayanos 2008).

The organization of this paper is as follows. In section 2, we present a graphical history of the inflation-indexed bond markets in the US and the UK, discussing bond supplies, the levels of yields, and the volatility and covariances with stocks of high-frequency movements in yields. In section 3, we ask what portion of the TIPS yield history can be explained by movements in short-term real interest rates, together with the expectations hypothesis of the term structure. This section revisits the VAR analysis of Campbell and Shiller (1996). In section 4, we discuss the risk characteristics of TIPS and estimate a model of TIPS pricing with time-varying systematic risk, a variant of Campbell, Sunderam, and Viceira (2009), to see how much of the yield history can be explained by changes in risk. In section 5, we discuss the unusual market conditions that prevailed in the fall of 2008 and the channels through which they influenced inflation-indexed bond yields. Section 6 draws implications for investors and policymakers. An Appendix available online (Campbell, Shiller, and Viceira 2009) presents technical details of our bond pricing model and of data construction.

2 The History of Inflation-Indexed Bond Markets

In this section we summarize graphically the history of two of the largest and best established inflation-indexed bond markets, the US TIPS market and the UK inflation-indexed gilt (UK government bond) market. We present a series of comparably formatted figures, first for the US (panel A of each figure) and then for the UK (panel B).

Figure 1A shows the growth of the outstanding supply of TIPS during the past ten years. From modest beginnings in 1997, the supply of TIPS grew to around 10% of the marketable debt of the US Treasury, and 3.5% of US GDP, in 2008. This growth has been fairly smooth, with a minor slowdown in 2001-02. Figure 1B shows a comparable history for the UK. From equally modest beginnings in 1982, inflation-

indexed gilts have grown rapidly to account for almost 30% of the British public debt, and 10% of GDP, in 2008. The growth in the inflation-indexed share of the public debt slowed down in 1990-97, and reversed in 2004-05, but otherwise proceeded at a rapid rate.

Figure 2A plots the yields on 10-year nominal and inflation-indexed US Treasury bonds over the period from January 1998 through March 2009. The figure shows a considerable decline in both nominal and real long-term interest rates since TIPS yields peaked early in the year 2000. Through 2007, the decline was roughly parallel, as inflation-indexed yields fell from slightly over 4% to slightly over 1%, while nominal yields fell from around 7% to 4%. Thus, this was a period in which both nominal and inflation-indexed bond yields were driven down by a large decline in long-term real interest rates. In 2008, however, nominal Treasury bond yields continued to decline, while inflation-indexed bond yields spiked up above 3% towards the end of the year.

Figure 2B shows a comparable history for the UK since the early 1990's. To facilitate comparison of the two plots, the beginning of the US sample period is marked with a vertical dashed line. The downward trend in inflation-indexed government bond yields is even more dramatic over this longer period. UK inflation-indexed gilts also experienced a dramatic yield spike in the fall of 2008.

Figure 3A plots the 10-year break-even inflation rate, the difference between 10-year nominal and inflation-indexed bond yields. The breakeven inflation rate was fairly volatile in the first few years of the TIPS market, then stabilized between 1.5% and 2.0% in the early years of this decade before creeping up to a stable level of about 2.5% from 2004 through 2007. In 2008, the breakeven inflation rate collapsed, reaching almost zero at the end of the year.

The figure also shows, for the early years of the sample, the subsequently realized 3-year inflation rate. After the first couple of years, in which there is little relation between breakeven and subsequently realized inflation, one can see that a slight decrease in breakeven inflation between 2000 and 2002, followed by a slow increase in breakeven inflation from 2002 to 2006, is matched by similar gradual changes in subsequently realized inflation. Although this is not a rigorous test of the rationality of the TIPS market—apart from anything else, the bonds are forecasting inflation over 10 years, not 3 years—it does suggest that inflation forecasts influence the relative pricing of TIPS and nominal Treasury bonds. We explore this issue in greater detail in the next section of the paper.

Figure 3B reports the breakeven inflation history for the UK. The figure shows a strong decline in breakeven inflation in the late 1990's, probably associated with the independence granted to the Bank of England by the newly elected Labour government in 1997, and a steady upward creep from 2003 to early 2008, followed by a collapse in 2008 comparable to that which has occurred in the US.

In Figure 4A we turn our attention to the short-run volatility of TIPS returns. Using daily nominal prices, with the appropriate correction for coupon payments, we calculate daily nominal return series for 10-year TIPS. The figure plots the annualized standard deviation of this series within a moving one-year window. For comparison, the figure also shows the corresponding annualized standard deviation for 10-year nominal Treasury bond returns, calculated from Bloomberg yield data using the assumption that nominal bonds trade at par.

The striking message of Figure 4A is that TIPS returns have become far more volatile in recent years. In the early years, until 2002, the short-run volatility of 10-year TIPS was only about half the short-run volatility of 10-year nominal Treasuries, but the two standard deviations converged between 2002 and 2004 and have been extremely similar since then. The annualized standard deviation of both bonds ranged between 5% and 8% until 2008, and then increased dramatically to about 13% during 2008.

Mechanically, two variables drive the volatility of TIPS returns. The most important is the volatility of TIPS yields, which has increased over time; in recent years it has been very similar to the volatility of nominal yields as breakeven inflation has stabilized. A second, amplifying factor is the duration of TIPS, which has increased as TIPS yields have declined.² The same two variables determine the very similar volatility patterns shown in Figure 4B for the UK.

Figure 5A plots the annualized standard deviation of 10-year breakeven inflation (a bond position long a 10-year nominal Treasury and short a 10-year TIPS). This standard deviation trended down from 6% in 1998 to about 1% in 2007, before spiking up above 13% in 2008. To the extent that breakeven inflation represents the long-term

²The duration of a bond is the weighted average time to payment of its cash flows, where the present values of cash flows are used as weights. Duration also equals the elasticity of a bond's price with respect to its gross yield. Coupon bonds have duration less than their maturity, and duration increases as yield falls. Since TIPS yields are lower than nominal yields, TIPS have greater duration for the same maturity, and hence a greater return volatility for the same yield volatility, but the differences in volatility explained by duration are quite small.

inflation expectations of market participants, these expectations stabilized during most of our sample period, but moved dramatically in 2008. Such a destabilization of inflation expectations should be a matter of serious concern to the Federal Reserve, although, as we discuss in section 5, institutional factors may have contributed to the movements in breakeven inflation during the market disruption of late 2008. Figure 5B shows that the Bank of England should be equally concerned by the recent destabilization of the yield spread between nominal and inflation-indexed gilts.

The figures also plot the correlations of daily inflation-indexed and nominal bond returns within a one-year moving window. Early in the period, the US correlation was quite low at about 0.2, but it increased to almost 0.9 by the middle of 2003 and stayed there until 2008. In the mid-2000's, TIPS behaved like nominal Treasuries and did not exhibit independent return variation. This coupling of TIPS and nominal Treasuries ended in 2008. The same patterns are visible in the UK data.

Although TIPS have been volatile assets, this does not necessarily imply that they should command large risk premia. According to rational asset pricing theory, risk premia should be driven by assets' covariances with the marginal utility of consumption rather than by their variances. One common proxy for marginal utility, used in the Capital Asset Pricing Model (CAPM), is the return on an aggregate equity index. In Figures 6A and 6B we plot the correlations of daily inflation-indexed bond returns, nominal government bond returns, and breakeven inflation returns (the difference between the first two series) with the daily returns on aggregate US and UK stock indexes, within our standard moving one-year window. Figures 7A and 7B repeat this exercise for betas (regression coefficients of daily bond returns and breakeven inflation onto the stock index).

All these figures tell a similar story. During the 2000's there has been considerable instability in the correlations between US and UK government bonds and stock returns, but these correlations have been predominantly negative, implying that government bonds can be used to hedge equity risk. To the extent that the CAPM describes risk premia across asset classes, government bonds should have predominantly negative rather than positive risk premia. The negative correlation is particularly striking for nominal government bonds, because breakeven inflation is positively correlated with stock returns, especially during 2002-03 and 2007-08. Campbell, Sunderam, and Viceira (2009) build a model in which a changing correlation between inflation and stock returns drives changes in the risk properties of nominal Treasury bonds. Their model assumes a constant equity market correlation for TIPS, and thus

cannot explain the correlation movements shown for TIPS in Figures 6A and 7A. In section 4 of this paper, we explore the determination of TIPS risk premia in greater detail.

3 Inflation-Indexed Yields and the Dynamics of Short-Term Real Interest Rates

To understand the movements of inflation-indexed bond yields, it is essential first to understand how changes in short-term real interest rates propagate along the real term structure. Declining yields for inflation-indexed bonds in the 2000's may not be particularly surprising given that short-term real interest rates have also been low in this decade.

Before TIPS were issued in 1997, Campbell and Shiller (1996) used a time-series model for the short-term real interest rate to create a hypothetical TIPS yield series under the assumption that the expectations theory of the term structure in log form, with zero log risk premia, describes inflation-indexed yields. (This does not require the assumption that the expectations theory describes nominal yields, a model that has often been rejected in US data.) In this section, we update Campbell and Shiller's analysis and ask how well the simple expectations theory describes the 12-year history of TIPS yields.

Campbell and Shiller estimated a VAR model in quarterly US data over the period 1953-1994. Their basic VAR included the ex post real return on a 3-month nominal Treasury bill, the nominal bill yield, and the lagged one-year inflation rate, with a single lag. They solved the VAR forward to create forecasts of future quarterly real interest rates at all horizons, and then aggregated the forecasts to generate the implied long-term inflation-indexed bond yield.

In Table 1A, we repeat this analysis for the period 1982-2008. The top panel reports the estimates of VAR coefficients, with standard errors in parentheses below. The bottom panel reports selected sample moments of the hypothetical VAR-implied 10-year TIPS yields, and for comparison the same moments of observed TIPS yields, over the period since TIPS were issued in 1997. The table delivers several interesting results.

First, hypothetical yields are considerably lower on average than observed yields, with a mean of 1.17% as compared with 2.68%. This implies that on average, investors demand a risk or liquidity premium for holding TIPS rather than nominal Treasuries. Second, hypothetical yields are more stable than observed yields, with a standard deviation of 0.36% as opposed to 0.94%. This reflects the fact that observed yields have declined more dramatically since 1997 than have hypothetical yields. Third, hypothetical and observed yields have a relatively high correlation of 0.70, even though no TIPS data were used to construct the hypothetical yields. Real interest rate movements do have an important effect on the TIPS market, and the VAR system is able to capture much of this effect.

Figure 8A shows these results in graphical form, plotting the history of the observed TIPS yield, the hypothetical VAR-implied TIPS yield, and the VAR estimate of the ex ante short-term real interest rate. The sharp decline in the real interest rate in 2001 and 2002 drives down the hypothetical TIPS yield, but the observed TIPS yield is more volatile and declines more strongly. The gap between the observed TIPS yield and the hypothetical yield shrinks fairly steadily over the sample period until the very end, when the 2008 spike in observed yields widens the gap again. These results suggest that when they were first issued, TIPS commanded a high risk or liquidity premium, which declined until 2008.

Table 1B and Figure 8B repeat these exercises for the UK. The hypothetical and observed yields have very similar means in the UK (2.64% and 2.67% respectively), but again the standard deviation is lower for hypothetical yields at 0.66% than for actual yields at 1.03%. The two yields have a high correlation of 0.79. Figure 8B shows that the VAR model captures much of the decline in inflation-indexed gilt yields since the early 1990's. It is able to do this because the estimated process for the UK ex ante real interest rate is highly persistent, so the decline in the real rate over the sample period translates almost one for one into a declining yield on long-term inflation-indexed gilts. However, for the same reason the model cannot account for variations in the yield spread between the short-term expected real interest rate and the long-term inflation-indexed gilt yield in the UK.

It is notable that the expectations hypothesis of the real term structure does not explain the decline in UK inflation-indexed gilt yields from 2005 through 2008. A change in UK accounting standards, FRS 17, may account for this. As Viceira (2003) and Vayanos and Vila (2007) explain, FRS 17 requires UK pension funds to mark their liabilities to market, using discount rates derived from government bonds.

The accounting standard was implemented, after some delay, in 2005, and it greatly increased the demand for inflation-indexed gilts from pension funds seeking to hedge their inflation-indexed liabilities.

4 The Systematic Risks of Inflation-Indexed Bonds

The yield history and VAR analysis presented in the previous two sections suggest that inflation-indexed bonds had low risk premia in the mid-2000's, but, in the US at least, had higher risk premia when they were first issued. In this section we use asset pricing theory to ask what fundamental properties of the macroeconomy might lead to high or low risk premia on inflation-indexed bonds. We first use the consumption-based asset pricing framework, and then present a less structured empirical analysis that relates bond risk premia to changing covariances of bonds with stocks.

4.1 Consumption-Based Pricing of Inflation-Indexed Bonds

A standard paradigm for consumption-based asset pricing assumes that a representative investor has Epstein-Zin (1989, 1991) preferences. This preference specification, a generalization of power utility, allows the coefficient of relative risk aversion γ and the elasticity of intertemporal substitution (EIS) ψ to be separate free parameters, whereas power utility restricts one to be the reciprocal of the other.

Under the additional assumption that asset returns and consumption are jointly lognormal and homoskedastic, the Epstein-Zin Euler equation implies that the risk premium on any asset i over the short-term safe asset is

$$RP_i \equiv E_t[r_{i,t+1}] - r_{f,t+1} + \frac{\sigma_i^2}{2} = \theta \frac{\sigma_{ic}}{\psi} + (1 - \theta)\sigma_{iw}. \quad (1)$$

The risk premium is defined to be the expected excess log return on the asset plus one-half its variance to correct for Jensen's Inequality. The preference parameter $\theta \equiv (1 - \gamma)/(1 - 1/\psi)$; in the power utility case, $\gamma = 1/\psi$ and $\theta = 1$. According to this formula, the risk premium on any asset is a weighted average of two conditional covariances, the consumption covariance σ_{ic} (scaled by the reciprocal of the EIS) which gets full weight in the power utility case, and the wealth covariance σ_{iw} . The risk premium is constant over time by the assumption of homoskedasticity.

It is tempting to treat the consumption covariance and wealth covariance as two separate quantities, but this ignores the fact that consumption and wealth are linked by the intertemporal budget constraint and by a time-series Euler equation. By using these additional equations, one can substitute either consumption (Campbell 1993) or wealth (Restoy and Weil 1998) out of the formula for the risk premium.

The first approach explains the risk premium using covariances with the current market return and with news about future market returns; this might be called “CAPM+”, as it generalizes the insight about risk that was first formalized in the CAPM. Campbell (1996) and Campbell and Vuolteenaho (2004) pursue this approach, which can also be regarded as an empirical version of Merton’s (1973) intertemporal CAPM.

The second approach explains the risk premium using covariances with current consumption growth and with news about future consumption growth; this might be called the “CCAPM+”, as it generalizes the insight about risk that is contained in the consumption-based CAPM with power utility. This approach has generated a large asset pricing literature in recent years (Bansal and Yaron 2004, Bansal, Khatchatrian, and Yaron 2005, Piazzesi and Schneider 2006, Bansal, Kiku, and Yaron 2007, Bansal, Dittmar, and Kiku 2008, Hansen, Heaton, and Li 2008). Some of this recent work adds heteroskedasticity to the simple homoskedastic model discussed here.

The CAPM+ approach delivers an approximate formula for the risk premium on any asset as

$$RP_i = \gamma\sigma_{iw} - (\gamma - 1)\sigma_{i,TIPS},$$

where σ_{iw} is the covariance of the unexpected return on asset i with the return on the aggregate wealth portfolio, and $\sigma_{i,TIPS}$ is the covariance with the return on an inflation-indexed perpetuity.

The intuition, which dates back to Merton (1973), is that conservative long-term investors value assets that deliver high returns at times when investment opportunities are poor. Such assets hedge investors against variation in the sustainable income stream that is delivered by a given amount of wealth. In a homoskedastic model, risk premia are constant and the relevant measure of long-run investment opportunities is the yield on an inflation-indexed bond. Thus, the covariance with the return on an inflation-indexed perpetuity captures the intertemporal hedging properties of an asset. In equilibrium, an asset that covaries strongly with an inflation-indexed perpetuity will offer a low return as the price of the desirable insurance it offers.

Applying this formula to the inflation-indexed perpetuity itself, we find that

$$RP_{TIPS} = \gamma \sigma_{TIPS,w} - (\gamma - 1) \sigma_{TIPS}^2.$$

The risk premium on a long-term inflation-indexed bond is increasing in its covariance with the wealth portfolio, as in the traditional CAPM, but decreasing in the variance of the bond return whenever the risk aversion of the representative agent is greater than one. Paradoxically, the insurance value of inflation-indexed bonds is higher when these bonds have high short-term volatility, because in this case they hedge important variability in investment opportunities. In a traditional model with a constant real interest rate, inflation-indexed bonds have constant yields; but in this case there is no intertemporal hedging to be done, and the traditional CAPM can be used to price all assets including inflation-indexed bonds.

The CCAPM+ approach can be written as

$$RP_i = \gamma \sigma_{ic} + \left(\gamma - \frac{1}{\psi} \right) \sigma_{ig}, \quad (2)$$

where σ_{ig} is the covariance of the unexpected return on asset i with revisions in expected future consumption growth \tilde{g}_{t+1} , defined by

$$\tilde{g}_{t+1} \equiv (\mathbf{E}_{t+1} - \mathbf{E}_t) \sum_{j=1}^{\infty} \rho^j \Delta c_{t+1+j}. \quad (3)$$

The risk premium on any asset is the coefficient of risk aversion γ times the covariance of that asset with consumption growth, plus $(\gamma - 1/\psi)$ times the covariance of the asset with revisions in expected future consumption growth. The second term is zero if $\gamma = 1/\psi$, the power utility case, or if consumption growth is unpredictable so that there are no revisions in expected future consumption growth. Evidence on the equity premium and the time-series behavior of real interest rates suggests that $\gamma > 1/\psi$. This implies that controlling for assets' contemporaneous consumption covariance, investors require a risk premium to hold assets that pay off when expected future consumption growth increases. Bansal and Yaron (2004) use the term “long-run risks” to emphasize this property of the model.

What does this model imply about the pricing of an inflation-indexed perpetuity? When expected real consumption growth increases by 1 percentage point,

the equilibrium real interest rate increases by $1/\psi$ percentage points, and thus the inflation-indexed perpetuity return is given by³

$$r_{TIPS,t+1} = -\frac{1}{\psi} \tilde{g}_{t+1}. \quad (4)$$

Combining (2) with (4), we can solve for the risk premium on the inflation-indexed perpetuity:

$$RP_{TIPS} = \gamma \left(-\frac{1}{\psi} \right) \sigma_{cg} + \left(\gamma - \frac{1}{\psi} \right) \left(-\frac{1}{\psi} \right) \sigma_g^2. \quad (5)$$

With power utility, only the first term in (5) is nonzero. This case is described by Campbell (1986). In a consumption-based asset pricing model with power utility, assets are risky if their returns covary positively with consumption growth. Since bond prices rise when interest rates fall, bonds are risky assets if interest rates fall in response to consumption growth. Because equilibrium real interest rates are positively related to expected future consumption growth, this is possible only if positive consumption shocks drive down expected future consumption growth, that is, if consumption growth is negatively autocorrelated. In an economy with temporary downturns in consumption, equilibrium real interest rates rise and TIPS prices fall in recessions, so investors require a risk premium to hold TIPS.

In the presence of persistent shocks to consumption growth, by contrast, consumption growth is positively autocorrelated. In this case recessions not only drive down current consumption but lead to prolonged periods of slow growth, driving down real interest rates. In such an economy the prices of long-term inflation-indexed bonds rise in recessions, making them desirable hedging assets with negative risk premia.

This paradigm suggests that the risk premium on TIPS will fall if investors become less concerned about temporary business-cycle shocks, and more concerned about shocks to the long-term consumption growth rate. It is possible that such a shift in investor beliefs did take place during the late 1990's and 2000's, as the Great Moderation mitigated concerns about business-cycle risk while long-term uncertainties about technological progress and climate change became more salient. Of course,

³A more careful derivation of this expression can be found in Campbell (2003), equation (34) on p.839.

the events of 2007-08 have brought business cycle risk to the fore again. The movements of inflation-indexed bond yields have been broadly consistent with changing risk perceptions of this sort.

The second term in (5) is also negative under the plausible assumption that $\gamma > 1/\psi$, and its sign does not depend on the persistence of the consumption process. However its magnitude does depend on the volatility of shocks to long-run expected consumption growth. Thus increasing uncertainty about long-run growth drives down inflation-indexed bond premia through this channel as well.

Overall, the Epstein-Zin paradigm suggests that inflation-indexed bonds should have low or even negative risk premia relative to short-term safe assets, consistent with the intuition that these bonds are the safe asset for long-term investors.

4.2 Bond Risk Premia and the Bond-Stock Covariance

The consumption-based analysis of the previous section delivers insights but also has weaknesses. The model assumes constant second moments and thus implies constant risk premia; it cannot be used to track changing variances, covariances, or risk premia in the inflation-indexed bond markets. While one could generalize the model to allow time-varying second moments, as in the long-run risks model of Bansal and Yaron (2004), the low frequency of consumption measurement makes it difficult to implement the model empirically. In this section, we follow a different approach, writing down a model of the stochastic discount factor (SDF) that allows us to relate the risk premia on inflation-indexed bonds to the covariance of these bonds with stock returns.

In order to capture the time-varying correlation of the returns on inflation-indexed bonds with stock returns, we propose a highly stylized term structure model in which the real interest rate is subject to conditionally heteroskedastic shocks. Conditional heteroskedasticity is driven by a state variable which captures time variation in aggregate macroeconomic uncertainty. We build our model in the spirit of Campbell, Sunderam, and Viceira (2009), which emphasizes the importance of changing macroeconomic conditions to understand time variation in systematic risk and in the correlation of returns on fundamental asset classes. Our model modifies their quadratic term structure model to allow for heteroskedastic shocks to the real rate.

We assume that the log of the real SDF, $m_{t+1} = \log M_{t+1}$, can be described by

$$-m_{t+1} = x_t + \frac{1}{2}\sigma_m^2 + \varepsilon_{m,t+1}, \quad (6)$$

where x_t follows a conditionally heteroskedastic AR(1) process,

$$x_{t+1} = \mu_x(1 - \phi_x) + \phi_x x_t + v_t \varepsilon_{x,t+1} + \varepsilon'_{x,t+1}, \quad (7)$$

and v_t follows a standard AR(1) process

$$v_{t+1} = \mu_v(1 - \phi_v) + \phi_v v_t + \varepsilon_{v,t+1}. \quad (8)$$

The shocks $\varepsilon_{m,t+1}$, $\varepsilon_{x,t+1}$, $\varepsilon'_{x,t+1}$, and $\varepsilon_{v,t+1}$ have zero means and are jointly normally distributed shocks with constant variance-covariance matrix. We assume that $\varepsilon'_{x,t+1}$ and $\varepsilon_{v,t+1}$ are orthogonal to each other and to the other shocks in the model. We adopt the notation σ_i^2 to describe the variance of shock ε_i , and σ_{ij} to describe the covariance between shock ε_i and shock ε_j . The conditional volatility of the log SDF (σ_m) describes the price of aggregate market risk or maximum Sharpe ratio in the economy, which we assume to be constant.⁴

In the Appendix, we show how to solve this model for the real term structure of interest rates. The state variable x_t is equal to the log short-term real interest rate, which follows an AR(1) process whose conditional variance is driven by the state variable v_t .

In a standard consumption-based power utility model of the sort we discussed in the previous subsection, v_t would capture time-variation in the dynamics of consumption growth. When v_t is close to zero, shocks to the real interest rate are uncorrelated with the stochastic discount factor; in a power utility model, this would imply that shocks to future consumption growth are uncorrelated with shocks to the current level of consumption. As v_t moves away from zero, the volatility of the real interest rate increases and its covariance with the SDF becomes more positive or more negative. In a power utility model, this corresponds to a covariance between consumption shocks

⁴CSV consider a much richer term structure model in which σ_m^2 is time varying. They note that in that case the process for the log real SDF admits an interpretation as a reduced model of structural models such as those of Bekaert, Engstrom and Grenadier (2005) and Campbell and Cochrane (1999) in which aggregate risk aversion is time-varying. CSV find that time-varying risk aversion plays only a limited role in explaining the observed variation in bond risk premia. For simplicity, we set σ_m^2 to be constant.

and future consumption growth that is either positive or negative, reflecting either momentum or mean-reversion in consumption. Broadly speaking we can interpret v_t as a measure of aggregate uncertainty about long-run growth in the economy. At times where uncertainty about future economic growth increases, real interest rates become more volatile.

Solving the model for the real term structure of interest rates, we find that the price of an n -period log inflation-indexed bond is linear in the short-term real interest rate x_t , with coefficient $B_{x,n}$, and quadratic in aggregate economic uncertainty v_t , with linear coefficient $B_{v,n}$ and quadratic coefficient $C_{v,n}$. An important property of this model is that bond risk premia are time varying. They are approximately linear in v_t , where the coefficient on v_t is proportional to σ_m^2 .

A time varying conditional covariance between the SDF and the real interest rate implies that the conditional covariance between real bonds and risky assets such as equities should also vary over time as a function of v_t . To see this, we now introduce equities into the model. To keep things simple, we assume that the unexpected log return on equities is given by

$$r_{e,t+1} - \mathbb{E}_t r_{e,t+1} = \beta_{em} \varepsilon_{m,t+1}. \quad (9)$$

This implies that the equity premium equals $\beta_{em} \sigma_m^2$, the conditional standard deviation of stock returns is $\beta_{em} \sigma_m$, and the Sharpe ratio on equities is σ_m . Equities deliver the maximum Sharpe ratio because they are perfectly correlated with the SDF. Thus we are imposing the restrictions of the traditional CAPM, ignoring the intertemporal hedging arguments given in the previous subsection.

The covariance between stocks and inflation-indexed bonds is given by

$$\text{Cov}_t(r_{e,t+1}, r_{n,t+1}) = B_{x,n-1} \beta_{em} \sigma_m x v_t, \quad (10)$$

which is proportional to v_t . This proportionality is also a reason why we consider two independent shocks to x_t . In the absence of a homoskedastic shock $\varepsilon'_{x,t}$ to x_t , our model would imply that the conditional volatility of the short real rate would be proportional to the covariance of stock returns with real bond returns. However, while the two moments appear to be correlated in the data, they are not perfectly correlated, still less proportional to one another.

We estimate this term structure model using the nonlinear Kalman Filter procedure described in CSV. To estimate the model we use data on zero-coupon inflation-

indexed bond yields from Gürkaynak, Sack and Wright (2008) for the period 1999-2008, and data on total returns on the value-weighted US stock market portfolio (inclusive of NYSE, NASDAQ and AMEX) from CRSP.⁵ Because the US Treasury does not issue TIPS with short maturities, and there are no continuous observations of yields on near-to-maturity TIPS, this dataset does not include short-term zero-coupon TIPS yields. To approximate the short-term real interest rate we use the ex ante short-term real interest rate implied by our VAR approach described in Section 3.

In our estimation we make several identifying and simplifying assumptions. First, we identify σ_m using the long-run average Sharpe ratio for US equities, which we set to 0.23 on a quarterly basis (equivalent to 0.46 on an annual basis). Second, we identify β_{em} as the sample standard deviation of equity returns in our sample period (0.094 per quarter, or 18.9% per year) divided by σ_m , for a value of 0.41. Third, we exactly identify x_t with the ex-ante short-term real interest rate estimated from the VAR model of the previous section, which we treat as observed, adjusted by a constant.. That is, we give the Kalman filter a measurement equation that equates the VAR-estimated short-term real rate to x_t with a free constant term but no measurement error. The inclusion of the constant term is intended to capture liquidity effects which lower the yields on Treasury bills relative to the longer-term real yield curve.

Fourth, because the shock $\varepsilon_{x,t+1}$ is always premultiplied by v_t , we normalize σ_x to one. Fifth, we assume that there is perfect correlation between the shock $\varepsilon_{x,t+1}$ and the shock $\varepsilon_{m,t+1}$ to the SDF; equivalently, we set σ_{mx} equal to 0.23. This delivers the largest possible time-variation in inflation-indexed bond risk premia, and thus maximizes the effect of changing risk on the TIPS yield curve. Sixth, we treat equation (10) as a measurement equation with no measurement error, where we replace the covariance on the left-hand side of the equation with the realized monthly covariance of returns on 10-year zero-coupon TIPS with return on stocks. We estimate the monthly realized covariance using daily observations on stock returns and on TIPS returns from the Gürkaynak-Sack-Wright dataset. Since β_{em} and σ_{mx} have been already exactly identified, this is equivalent to identifying the process v_t with a scaled version of the covariance of returns on TIPS and stocks.

⁵Gürkaynak, Sack and Wright estimate zero-coupon TIPS yields by fitting a flexible functional form, a generalization of Nelson and Siegel (1987) suggested by Svensson (1994), to the instantaneous forward rates implied by off-the-run TIPS yields. From fitted forward rates it is straightforward to obtain zero coupon yields.

We include one final measurement equation for the 10-year zero-coupon TIPS yield using the model's solution for this yield and allowing for measurement error. The identifying assumptions we have made imply that we are exactly identifying x_t with the observed short-term real rate, v_t with the realized covariance of returns on TIPS and stocks, and the log SDF with stock returns. Thus our estimation procedure in effect generates hypothetical TIPS yields from these processes, and compares them with observed TIPS yields.

Table 2 reports the parameter estimates from our full model, in the middle column, and two restricted models. The left column drops the measurement equation for the realized stock-bond covariance and assumes that the stock-bond covariance is constant, hence that TIPS have constant risk premia, as in the VAR model of section 3. The right hand column generates the largest possible effects of time-varying risk premia on TIPS yields by increasing the persistence of the covariance state variable v_t from the freely estimated value of 0.77, which implies an 8-month half-life for covariance movements, to the largest permissible value of one.

Figure 9 shows how these three variants of our basic model fit the history of the 10-year TIPS yield. The actual TIPS yield is the thick solid line in the figure. The dotted and thin solid lines are the freely estimated model of changing risk and the restricted model with a constant bond-stock covariance. The fact that these lines are almost on top of one another, diverging only slightly in periods such as 2003 and 2008 when the realized bond-stock covariance was unusually negative, tells us that changing TIPS risk is not persistent enough to have a large effect on TIPS yields. Only when we impose a unit root on the process for the bond-stock covariance, illustrated with a dashed line in the figure, do we obtain large effects of changing risk. This model implies that TIPS yields should have fallen more dramatically than they did in 2002-03, and again in 2007, when the covariance of TIPS with stocks turned negative. The model does capture TIPS movements in the first half of 2008, but dramatically fails to capture the spike in TIPS yields in the second half of 2008.

Overall, this exploration of changing risk, as captured by the changing realized covariance of TIPS returns and aggregate stock returns, suggests that risk variations play only a supporting role in the determination of TIPS yields. The major problem with a risk-based explanation for movements in the inflation-indexed yield curve is that the covariance of TIPS and stocks has moved in a transitory fashion, and thus should not have had a large effect on TIPS yields unless investors were expecting more persistent variation and were surprised by an unusual sequence of temporary

changes in risk.

These results contrast with those reported by CSV, who find that persistent movements in the covariance between inflation and stock returns have had a powerful influence on the nominal US Treasury yield curve. CSV find that US inflation was negatively correlated with stock returns in the late 1970's and early 1980's, when the major downside risk for investors was stagflation; it has been positively correlated with stock returns in the 2000's, when investors have been more concerned about deflation. As a result, CSV argue that the inflation risk premium was positive in the 1970's and 1980's but has been negative in the 2000's, implying even lower expected returns on nominal Treasury bonds than on TIPS. The movements in inflation risk identified by CSV are persistent enough to have important effects on the shape of the nominal US Treasury yield curve, reducing its slope and concavity relative to the shapes that were typical in the 1970's and 1980's. Figure 6A in this paper illustrates the positive correlation of US inflation and stock returns during the 2000's, while Figure 6B shows that this correlation has changed sign in the UK since the early 1990's.

5 The Crisis of 2008 and Institutional Influences on TIPS Yields

In 2008, as the subprime crisis intensified, the TIPS yield became highly volatile, and appeared suddenly disconnected from the yield on nominal Treasuries. At the beginning of 2008 the 30-year TIPS yield fell to extremely low levels, as low as 1.66% on January 23, 2008. Shorter maturity TIPS showed even lower yields, and in the summer of 2008 some of these showed yields below minus 0.5%, reminding market participants that zero is not the lower bound for inflation-indexed bond yields. Then, in the fall of 2008 there was an unprecedented and short-lived spike in TIPS yields, peaking at the end of October 2008 when the 30-year TIPS yield reached 3.44%.

These extraordinary short-run movements in TIPS yields are mirrored in the ten-year TIPS yield shown in Figure 2A. The extremely low TIPS yield in early 2008 was given a convenient explanation by some market observers, that people were panicked by the apparent heightened risks in financial markets brought in to the subprime crisis, and would buy safety at just about any price. But, if this is the explanation

of that phenomenon, we are left with the mystery of a massive surge in TIPS yield later in that year. The leap upwards in TIPS yields in the fall of 2008 was puzzling since it was not shared by nominal bond yields, and so it marked a massive drop in the breakeven inflation rate, seen in Figure 3A. The UK market behaved in a similar fashion, as can be seen from Figures 2B and 3B.

The anomalous sudden jump in inflation-indexed bond yields came as a total surprise to market participants. Indeed, just as the sudden jump occurred in October 2008 some observers were saying that because inflation expectations had become extremely stable, TIPS and nominal Treasury bonds were virtually interchangeable. Brière and Signori (2008) concluded in a paper published in October 2008 that “Although diversification was a valuable reason for introducing IL bonds in a global portfolio before 2003, this is no longer the case.” The extent of this surprise suggests that the TIPS yield, and its decoupling from the nominal Treasury yields, had something to do with the systemic nature of the crisis that beset US financial institutions in 2008.

Indeed, the sharp peak in the TIPS yield and the companion steep drop in the breakeven inflation rate occurred shortly after an event that some observers blame for the anomalous behavior of TIPS yields. This was the bankruptcy of the investment bank Lehman Brothers, announced on September 15, 2008. The unfolding of the Lehman Brothers bankruptcy proceedings also took place over the same interval of time over which the inflation-indexed bond yield made its spectacular leap upwards.

Lehman’s bankruptcy was an important event, the first bankruptcy of a major investment bank since the Drexel Burnham Lambert bankruptcy in 1990. That is not to say that other investment banks did not get into trouble, especially during the subprime crisis. But, the government had always stepped in to allay fears. Bear Stearns was sold to commercial bank J.P. Morgan in March 2008 in a deal arranged and financed by the government. Bank of America announced its purchase of Merrill Lynch on September 14, 2008, again with government financial support. The government decided to let Lehman fail, and so it is possible that this event was indicative of future government policy that might spell major changes in the economy.

One conceivable interpretation for the events that followed the Lehman bankruptcy announcement is that the bankruptcy was seen by the market as a macroeconomic indicator, indicating that the economy would be suddenly weaker. This could imply a deterioration in the government’s fiscal position, justifying an increase in expected future real interest rates and therefore the long-term real yield on US Treas-

sury debt, and a decline in inflation expectations, explaining the drop in breakeven inflation.

However, many observers doubt that we can really expect such a radical change in real-rate and inflation expectations from the macroeconomic impact of just this one bankruptcy. At one point in 2008 the breakeven 7-year inflation rate reached minus 1.5%. According to a paper by Hu and Worah (2009), bond traders at PIMCO, “The market did not believe that it was possible to realize that kind of real rate or sustained deflation.”

Another interpretation of this shift is that there was a shift in the risk premium for inflation-indexed bonds. In terms of our analysis above, it could be a change in the covariance of TIPS returns with consumption or wealth. But, such a view sounds even less plausible than the view that the Lehman effect worked through inflation expectations. We have seen that the observed fluctuations in the covariances of TIPS returns with other variables are hard to rationalize even after the fact, and so it is hard to see why the market would have made a major adjustment in this covariance.

Hu and Worah conclude instead that “the extremes in valuation were due to a potent combination of technical factors.... Lehman owned Tips as part of repo trades or posted Tips as counterparty collateral. Once Lehman declared bankruptcy, both the court and its counterparty needed to sell these Tips for cash.” The traders at PIMCO saw then a flood of TIPS on the market. There appeared to be few buyers for these. Distressed market makers were not willing to risk taking positions in these TIPS; their distress was marked by a crisis-induced sudden and catastrophic widening, by October of 2008, in TIPS bid-asked spreads. The situation was exacerbated by the fact that some TIPS funds had commodity overlay strategies that forced them to sell TIPS because of the fall at that time in commodity prices. Moreover, institutional money managers had to confront a sudden loss of client interest in relative value trades, trades that might have exploited the abnormally low breakeven inflation.

5.1 Inflation Derivatives Markets in the Fall of 2008

An important clue about the events of fall 2008 is provided by the diverging behavior of breakeven inflation rates in the TIPS cash market and breakeven inflation rates implied by zero-coupon inflation swaps during the months following the Lehman bankruptcy.

Zero-coupon inflation swaps are derivatives contracts where one of the parties pays the other cumulative CPI inflation over the term of the contract at maturity, in exchange for a predetermined fixed rate. This rate is known as the “synthetic” breakeven inflation rate because, if inflation grew at this fixed rate over the life of the contract, the net payment on the contract at maturity would be equal to zero. As with the “cash” breakeven inflation rate implied by TIPS and nominal Treasury bonds, this rate reflects both expected inflation over the relevant period as well as an inflation risk premium.

Figure 10 plots the cash inflation breakeven rate implied by off-the-run TIPS and nominal Treasury bonds maturing on July 2017 and the synthetic inflation breakeven rate for the 10-year zero-coupon inflation swap for the time period between July 2007 and April 2009. The figure also plots the TIPS asset swap spread—explained below. The figure shows that the two breakeven rates track each other very closely up to mid-September 2008, with the synthetic inflation breakeven rate being about 35-40 basis points larger than the cash breakeven inflation rate on average.

This difference in breakeven rates is typical under normal market conditions. According to analysts, it reflects among other things the cost of manufacturing pure inflation protection in the US. Most market participants supplying inflation protection in the US inflation swap market are levered investors such as hedge funds and banks proprietary trading desks. These investors typically hedge their inflation swap positions by simultaneously taking long positions in TIPS and short positions in nominal Treasuries in the asset swap market. A buying position in an asset swap is functionally similar to a levered position in a bond. In an asset swap, one party pays the cash flows on a specific bond, and receives in exchange LIBOR plus a spread known as the asset swap spread. Typically this spread is negative and its absolute magnitude is larger for nominal Treasuries than for TIPS. Thus a levered investor paying inflation—i.e. selling inflation protection—in an inflation swap faces a positive financing cost derived from his long-short TIPS-nominal Treasury position.

Figure 10 shows that starting in mid-September 2008, cash breakeven rates fell dramatically while synthetic breakeven rates did not fall nearly as much, while at the same time TIPS asset swap spreads increased from their normal levels of about -35 basis points to about +100 basis points. Although not shown in the figure, nominal Treasury asset swap spreads remained at their usual levels. That is, financing long positions in TIPS became extremely expensive relative to historical levels just as their cash prices fell abruptly.

There is no reason why declining inflation expectations should directly affect the cost of financing long positions in TIPS relative to nominal Treasuries. These two simultaneous changes suggest instead that we may have witnessed an episode of intense selling in the cash market with insufficient demand to absorb those sales—as described by Hu and Worah—and simultaneously another shortage of capital to finance levered positions in markets other than nominal Treasuries; that is, we may have witnessed a “liquidity” episode.

Under this interpretation, in the fall of 2008 the synthetic breakeven inflation rate was a better proxy for inflation expectations in the marketplace than the cash breakeven inflation rate, despite the fact that in normal times the inflation swap market is considerably less liquid than the cash TIPS market. The synthetic breakeven inflation rate declined from about 3% per year to about 1.5% at the trough. This long-run inflation expectation is perhaps more plausible than the 0% 10-year inflation expectations reflected in the cash market for the off-the-run 2017 bonds.

Interestingly, cash breakeven inflation rates also diverged between on-the-run and off-the-run TIPS with similar maturities during this period. The online Appendix shows that breakeven rates based on newly issued, or on-the-run, TIPS were lower than those based on off-the-run TIPS. This divergence reflected another feature of TIPS which causes cash breakeven inflation rates calculated from on-the-run TIPS to be poor proxies for inflation expectations in the face of deflation risk. Contractually TIPS holders have the right to redeem their bonds at maturity for the maximum of either par value at issuance or that value plus accrued inflation during the life of the bond. Thus, when there is a risk of deflation after a period of inflation, new TIPS issues offer better deflation protection than old ones. Accordingly, on-the-run TIPS should be more expensive and thus their real yields lower than those of off-the-run TIPS. Breakeven inflation rates derived from on-the-run TIPS must be adjusted upwards for the deflation-protection premium to arrive at a measure of inflation expectations.

We view the experience with TIPS yields after the Lehman bankruptcy as the sign of a highly abnormal market situation, where liquidity problems suddenly created severe financial anomalies. This may seem to imply that we can take the recent episode as unrepresentative, and ignore the observations from these dates. And yet, investors in TIPS who would like to regard them as the safest long-term investments must consider the extraordinary short-term volatility that such events have given their yields.

6 The Uses of Inflation-Indexed Bonds

6.1 Implications for Investors

What lessons should investors learn from the history of TIPS and inflation-indexed gilt yields? The basic case for inflation-indexed bonds, stated by Campbell and Shiller (1996) and further developed by Brennan and Xia (2002), Campbell and Viceira (2001, 2002), Campbell, Chan, and Viceira (2003), and Wachter (2003), is that these bonds are the safe asset for long-term investors. An inflation-indexed perpetuity delivers a known stream of real spending power to an infinite-lived investor, and a zero-coupon inflation-indexed bond delivers a known real payment in the distant future to an investor who values wealth at that single horizon. This argument does not make any assumption about the time-series variation in yields, and so it is not invalidated by the gradual long-term decline in inflation-indexed bond yields since the 1990's, the mysterious medium-run variations in TIPS yields relative to short-term real interest rates, the spike in yields in the fall of 2008, or the high daily volatility of TIPS returns.

There are, however, two circumstances in which other assets can substitute for inflation-indexed bonds by providing long-term safe returns. First, if the breakeven inflation rate is constant, as will be the case when the central bank achieves perfect anti-inflationary credibility, then nominal bonds are perfect substitutes for inflation-indexed bonds and the conventional Treasury or gilt market can be used by conservative long-term investors. For a time in the mid-2000's, it looked as if we were approaching this nirvana of central bankers, but the events of 2008 dramatically destabilized inflation expectations and reaffirmed the distinction between real and nominal bonds.

Second, if the ex ante real interest rate is constant, as was famously asserted by Fama (1975), then long-term investors can roll over short-term Treasury bills to achieve almost perfectly certain long-term real returns. Because inflation uncertainty is minimal over a month or a quarter, Treasury bills expose investors to minimal inflation risk. In general, they do expose investors to the risk of persistent variation in the real interest rate, but this risk is absent if the real interest rate is constant over time.

Investors can tell whether this happy circumstance prevails by forecasting realized

real returns on Treasury bills and measuring the movements of their forecasts, as we did in Figures 8A and 8B, or more simply by measuring the volatility of inflation-indexed bond returns. If inflation-indexed bonds have yields that are almost constant and returns with almost no volatility, then Treasury bills are likely to be good substitutes.⁶ Seen from this point of view, the high daily volatility of inflation-indexed bond returns illustrated in Figures 4A and 4B, far from being a drawback, demonstrates the value of inflation-indexed bonds for conservative long-term investors.

A simple quantitative measure of the usefulness of inflation-indexed bonds is the reduction in long-run portfolio standard deviation that these bonds permit. We can estimate this reduction by calculating the long-run standard deviation of a portfolio of *other* assets chosen to minimize long-run risk. This is the smallest risk that long-run investors can achieve if inflation-indexed bonds are unavailable. Once inflation-indexed bonds become available, the minimum long-run risk portfolio consists entirely of these bonds and has zero long-run risk. Thus, the minimized long-run standard deviation of a portfolio of other assets measures the risk reduction that inflation-indexed bonds make possible.⁷

Figures 11 and 12 plot the time series of the annualized 10-year standard deviation of the real return on the 10-year global minimum variance (GMV) portfolio of US stocks, nominal 5-year Treasury bonds, and 3-month Treasury bills. To derive the composition of this portfolio and its volatility at each horizon we use the long-horizon mean-variance approach described in Campbell and Viceira (2005) and its companion technical guide (Campbell and Viceira 2004). We estimate a VAR(1) system for the ex-post real return on T-bills, and the excess log return on stocks and nominal bonds that also includes variables known to forecast bond and equity risk premia. Specifically, we include in the VAR system as forecasting variables the log dividend-price ratio, the spread between the yield on the 5-year bond and the yield on Treasury bills, and the yield on Treasury bills in addition to lagged returns and real interest

⁶Strictly speaking, this argument assumes that real yields are described by the expectations hypothesis of the term structure, so that constant short-term real interest rates imply constant long-term real yields. Volatile risk or liquidity premia on inflation-indexed bonds could make their yields volatile even if short-term real interest rates are constant. However, it is quite unlikely that time-variation in risk or liquidity premia would stabilize the yields on inflation-indexed bonds in an environment of time-varying real interest rates.

⁷As an alternative approach, Campbell, Chan, and Viceira (2003) calculate the utility of an infinite-lived investor who has access to stocks, nominal bonds, and bills, and the utility gain when this investor also can hold an inflation-indexed perpetuity. We do not update this more complex calculation here.

rates. From this system we extract the conditional variance-covariance of 10-year returns using the formulae in Campbell and Viceira (2004), and find the portfolio that minimizes this variance.

Instead of estimating a single VAR system for our entire quarterly sample, 1953:1-2008:4, we estimate two VAR systems, one for the period 1953:1-1972:4, and another for the period 1973:1-2008:4. We split the sample this way because we are concerned that the process for inflation and the real interest rate might have changed during this period. The conditional long-horizon moments of returns also depend on the quarterly variance-covariance matrix of innovations, which we estimate using three-year windows of quarterly data. Within each window and VAR sample period, we combine the variance-covariance matrix with the full-sample estimate of the slope coefficients to compute the 10-year GMV portfolio and its annualized volatility.

The annualized 10-year standard deviation of the 10-year GMV portfolio is fairly low in the 1960's at around 1% per year. This is the period that led Fama (1975) to assert that the ex ante real interest rate is constant over time. Starting in the 1970's, however, persistent movements in the real interest rate cause the standard deviation to rise rapidly to about 4% per year. The standard deviation drops back to about 2% in the mid-1990's, but by 2008 is once again at a historical high of 4%. These numbers imply that inflation-indexed bonds substantially reduce risk for long-term investors.

In Figure 11, we compare the estimated standard deviation of the GMV portfolio (right axis scale) with the annualized daily standard deviations of TIPS and inflation-indexed gilts (left axis scale) over the period where these bonds exist. In Figure 12, we compare the same GMV standard deviation with the estimated standard deviation of hypothetical TIPS returns, constructed from the VAR system using the method of Campbell and Shiller (1996) and section 3 of this paper, which assumes the log expectations hypothesis for inflation-indexed bonds. Both comparisons show that historically, the minimum long-run risk that can be achieved using other assets has been high when short-term TIPS returns have been volatile. In other words, inflation-indexed bonds are particularly good long-run risk reducers whenever their short-run risk is high. Such a result may seem paradoxical, but it follows directly from the fact that inflation-indexed bonds are needed for long-term safety when real interest rates vary persistently over time.⁸

⁸This point is related to the asset pricing result discussed in section 4.1, that controlling for the stock market covariance of inflation-indexed bonds, the equilibrium risk premium on these bonds

Inflation-indexed bonds also play an important role for institutional investors who need to hedge long-term real liabilities. Pension funds and insurance companies with multi-year commitments should use inflation-indexed bonds to neutralize the swings in the present value of their long-dated liabilities that will be caused by changes in long-term real interest rates. Of course, these swings only become apparent to institutional investors when they discount real liabilities using market real rates, as the UK has required in recent years. The resulting institutional demand for inflation-indexed gilts seems to have been an important factor driving down the yields on these bonds (Viceira 2003, Vayanos and Vila 2007).

The total demand of long-term investors for inflation-indexed bonds will depend not only on their risk properties, but also on their expected returns relative to other available investments and the risk tolerance of investors. An aggressive long-term investor might wish to short inflation-indexed bonds and invest the proceeds in equities, since stocks have only very rarely underperformed bonds over three or more decades in US and UK data. In 2008 it was reported that Clare College, Cambridge was planning to undertake such a strategy. However, Campbell, Chan, and Viceira (2003) estimated positive long-term demands for inflation-indexed bonds by long-term investors who also have the ability to borrow short-term or to issue long-term nominal bonds.

Long-term inflation-indexed bonds may be of interest to some short-term investors. Given their high short-run volatility, however, short-term investors will only wish to hold inflation-indexed bonds if they expect to receive high excess returns over Treasury bills (as might reasonably have been the case in 1999-2000, or during the yield spike of fall 2008), or if they hold other assets, such as stocks, whose returns can be hedged by an inflation-indexed bond position. We have shown evidence that TIPS and inflation-indexed gilts have hedged stock returns during the downturns of the early 2000's and the late 2000's, and this should make them attractive to short-term equity investors.

The illiquidity of inflation-indexed bonds is often mentioned as a disadvantage of the asset class. It is important to note, however, that inflation-indexed bonds are only illiquid relative to nominal government bonds which, along with foreign exchange, are the most liquid financial assets. Relative to almost any other long-term investment, inflation-indexed government bonds are extremely cheap to trade. In addition, long-

is declining in their variance when assets are priced by a conservative infinite-lived representative investor.

term buy-and-hold investors should care very little about transactions costs since they will rarely need to turn over their bond positions.

6.2 Implications for Policymakers

In managing the public debt, the US Treasury seeks to minimize the average cost of debt issue while paying due regard to risk, including refinancing risk. It is commonly thought that short-term Treasury bills are less expensive than long-term debt but that exclusive reliance on bills would impose an unacceptable refinancing risk as bills must frequently be rolled over.

In the period since TIPS were issued in 1997, they have proven to be an expensive form of debt *ex post*, because of the unexpected decline in real interest rates from the 1990's through early 2008. However, our analysis implies that the cost of TIPS should be lower than that of Treasury bills *ex ante* because TIPS offer investors desirable insurance against future variation in real interest rates. This is the relevant consideration going forward, as Roush, Dudley, and Steinberg Ezer (2008) emphasize, so governments should not be deterred from issuing inflation-indexed bonds by the high realized returns on their past inflation-indexed bond issues.

In the current environment, with inflation positively correlated with stock prices, the inflation risk premium in nominal Treasury bonds is likely negative. This implies that long-term nominal debt should be even cheaper for the Treasury than TIPS. However, the correlation between inflation and stock prices has changed sign in the past (Campbell, Sunderam, and Viceira 2009), and it may easily do so again in the future.

Several other considerations also suggest that inflation-indexed bonds are a valuable form of public debt. First, to the extent that particular forms of debt have different investment clienteles with downward-sloping demand curves for bonds, it is desirable to diversify across different forms to find the largest possible market for government debt (Greenwood and Vayanos 2008, Vayanos and Vila 2007).

Second, inflation-indexed bonds can be used to draw inferences about bond investors' inflation expectations, and such information is extremely valuable for monetary policymakers. It is true that market disruptions, such as those that occurred in the fall of 2008, complicate the measurement of inflation expectations, but our analy-

sis shows that it is possible to derive meaningful information even in these extreme conditions.

Finally, inflation-indexed bonds provide a safe real asset for long-term investors, and promote public understanding of inflation. These public benefits should be taken into account by fiscal authorities as part of their broader mission to improve the functioning of their economies.

References

- Backus, David and Stanley Zin. 1994. "Reverse Engineering the Yield Curve." NBER Working Paper No. 4676.
- Bansal, Ravi, and Amir Yaron. 2004. "Risks for the Long Run." *Journal of Finance* 59:1481–1509.
- Bansal, Ravi, Robert Dittmar, and Dana Kiku. 2008. "Cointegration and Consumption Risks in Asset Returns." Forthcoming *Review of Financial Studies*.
- Bansal, Ravi, Varoujan Khatchatrian, and Amir Yaron. 2005. "Interpretable Asset Markets?" *European Economic Review* 49:531–560.
- Bansal, Ravi, Dana Kiku and Amir Yaron. 2007. "Risks for the Long Run: Estimation and Inference." Unpublished paper, Duke University and University of Pennsylvania.
- Barr, David G., and John Y. Campbell. 1997. "Inflation, Real Interest Rates, and the Bond Market: A Study of UK Nominal and Index-Linked Government Bond Prices." *Journal of Monetary Economics* 39:361–383.
- Beechey, Meredith J., and Jonathan H. Wright. 2009. "The High-Frequency Impact of News on Long-Term Yields and Forward Rates: Is It Real?" Sveriges Riksbank and Johns Hopkins University.
- Bekaert, Geert, Eric Engstrom, and Steve Grenadier. 2005. "Stock and Bond Returns with Moody Investors." Unpublished paper, Columbia University, University of Michigan, and Stanford University.
- Brennan, Michael J., and Yihong Xia. 2002. "Dynamic Asset Allocation Under Inflation." *Journal of Finance* 57:1201–1238.
- Brière, Marie, and Ombretta Signori. 2008. "Do Inflation-Linked Bonds Still Diversify?" *European Financial Management*, 15 October.
- Campbell, John Y. 1986. "Bond and Stock Returns in a Simple Exchange Model." *Quarterly Journal of Economics* 101: 785–804.
- Campbell, John Y. 1993. "Intertemporal Asset Pricing without Consumption Data." *American Economic Review* 83:487–512.

- Campbell, John Y. 1996. "Understanding Risk and Return." *Journal of Political Economy* 104:298–345.
- Campbell, John Y. 2003. "Consumption-Based Asset Pricing." Chapter 13 in *Handbook of the Economics of Finance Vol. IB*, edited by George Constantinides, Milton Harris, and Rene Stulz. North-Holland.
- Campbell, John Y., and John H. Cochrane. 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy* 107:205–251.
- Campbell, John Y., and Robert J. Shiller. 1991. "Yield Spreads and Interest Rate Movements: A Bird's Eye View." *Review of Economic Studies* 58:495–514.
- Campbell, John Y., and Robert J. Shiller. 1996. "A Scorecard for Indexed Government Debt." In *National Bureau of Economic Research Macroeconomics Annual 1996*, edited by Ben S. Bernanke and Julio Rotemberg. MIT Press.
- Campbell, John Y., Robert J. Shiller, and Luis M. Viceira. 2009. "Understanding Inflation-Indexed Bond Markets: Appendix." Available online at <http://kuznets.fas.harvard.edu/~campbell/papers.html>.
- Campbell, John Y., Adi Sunderam, and Luis M. Viceira. 2009. "Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds." NBER Working Paper No. 14701.
- Campbell, John Y., and Luis M. Viceira. 2001. "Who Should Buy Long-Term Bonds?" *American Economic Review* 91:99–127.
- Campbell, John Y., and Luis M. Viceira. 2002. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press.
- Campbell, John Y., and Luis M. Viceira. 2004. "Long-Horizon Mean-Variance Analysis: A User Guide". Available online.
- Campbell, John Y., and Luis M. Viceira. 2005. "The Term Structure of the Risk-Return Tradeoff." *Financial Analysts Journal* 61(1):34–44.
- Campbell, John Y., and Tuomo O. Vuolteenaho. 2004. "Bad Beta, Good Beta." *American Economic Review* 94:1249–1275.

- Christensen, Jens H.E., Jose A. Lopez, and Glenn D. Rudebusch. 2008. “Inflation Expectations and Risk Premiums in an Arbitrage-Free Model of Nominal and Real Bond Yields.” Federal Reserve Bank of San Francisco.
- D’Amico, Stefania, Don H. Kim, and Min Wei. 2008. “Tips from TIPS: The Informational Content of Treasury Inflation-Protected Security Prices.” Federal Reserve Board of Governors.
- Epstein, Lawrence, and Stanley Zin. 1989. “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework.” *Econometrica* 57:937–69.
- Epstein, Lawrence, and Stanley Zin. 1991. “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Investigation.” *Journal of Political Economy* 99:263–286.
- Fabozzi, Frank J. 2005. *The Handbook of Fixed Income Securities*. McGraw-Hill.
- Fabozzi, Frank J., and Steven V. Mann. 2001. *Introduction to Fixed Income Analytics*. Frank J. Fabozzi Associates.
- Fama, Eugene F. 1975. “Short-Term Interest Rates as Predictors of Inflation.” *American Economic Review* 65:269–282.
- Fleming, Michael J., and Neel Krishnan. 2009. “The Microstructure of the TIPS Market.” Federal Reserve Bank of New York.
- Gollier, Christian. 2005. “The Consumption-Based Determinants of the Term Structure of Discount Rates.” University of Toulouse.
- Greenwood, Robin, and Dimitri Vayanos. 2008. “Bond Supply and Excess Bond Returns.” Harvard Business School and London School of Economics.
- Grishchenko, Olesya V., and Jing-zhi Huang. 2008. “Inflation Risk Premium: Evidence from the TIPS Market.” Pennsylvania State University.
- Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright. 2006. “The U.S. Treasury Yield Curve: 1961 to the Present.” FEDS working paper 2006-28, <http://www.federalreserve.gov/pubs/feds/2006/200628>.

- Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright. 2008. “The TIPS Yield Curve and Inflation Compensation.” Bilkent University, Macroeconomic Advisers, LLC, and Johns Hopkins University.
- Hansen, Lars Peter, John C. Heaton, and Nan Li. 2008. “Consumption Strikes Back? Measuring Long-Run Risk.” *Journal of Political Economy* 116:260–302.
- Haubrich, Joseph, George Pennacchi, and Peter Ritchken. 2008. “Estimating Real and Nominal Term Structures using Treasury Yields, Inflation, Inflation Forecasts, and Inflation Swap Rates.” Federal Reserve Bank of Cleveland, University of Illinois, and Case Western Reserve University.
- Hu, Gang, and Mihir Worah. 2009. “Why Tips Real Yields moved significantly higher after the Lehman Bankruptcy.” PIMCO, Newport Beach, CA .
- McCulloch, J. Huston, and H. Kwon. 1993. “US Term Structure Data, 1947–1991.” Working Paper 93-6, Ohio State University.
- Merton, Robert C. 1973. “An Intertemporal Capital Asset Pricing Model.” *Econometrica* 41:867–87.
- Nelson, Charles R., and Andrew F. Siegel. 1987. “Parsimonious Modeling of Yield Curves.” *Journal of Business* 60:473–489.
- Piazzesi, Monika, and Martin Schneider. 2006. “Equilibrium Yield Curves.” In *NBER Macroeconomics Annual 2006*, edited by Daron Acemoglu, Kenneth Rogoff, and Michael Woodford. MIT Press.
- Restoy, Fernando, and Philippe Weil. 1998. “Approximate Equilibrium Asset Prices.” NBER Working Paper No. 6611.
- Roush, Jennifer, William Dudley, and Michelle Steinberg Ezer. 2008. “The Case for TIPS: An Examination of the Costs and Benefits.” Staff Report 353, Federal Reserve Bank of New York.
- Svensson, Lars E.O. 1994. “Estimating and Interpreting Forward Interest Rates: Sweden 1992–1994.” NBER Working Paper No. 4871.
- Vayanos, Dimitri, and Jean-Luc Vila. 2007. “A Preferred-Habitat Model of the Term Structure of Interest Rates.” Unpublished paper, London School of Economics.

- Viceira, Luis M. 2003. “Pension Policy at The Boots Company PLC.” Case N2-203-105, Harvard Business School.
- Viceira, Luis M. 2007. “Bond Risk, Bond Return Volatility, and the Term Structure of Interest Rates.” Harvard Business School.
- Wachter, Jessica A. 2003. “Risk Aversion and Allocation to Long-Term Bonds.” *Journal of Economic Theory* 112:325–333.
- Wachter, Jessica A. 2006. “A Consumption-Based Model of the Term Structure of Interest Rates.” *Journal of Financial Economics* 79:365–399.

Table 1A
US VAR Estimation Results

Dependent Variable	Real T-bill rate	Nominal T-bill rate	Inflation	R ²
VAR Estimation Results				
Real T-bill return	-0.02 (0.12)	0.02 (0.03)	-0.19 (0.11)	0.25
Nominal T-bill yield	0.59 (0.18)	0.94 (0.04)	0.54 (0.18)	0.91
Inflation	0.09 (0.09)	-0.04 (0.02)	0.58 (0.08)	0.59
Const.	-0.005 (0.002)	-0.002 (0.0005)	0.007 (0.002)	
Observed and hypothetical moments of 10 year real bonds				
	Observed	Hypothetical		
Mean	2.66	1.04		
Standard Deviation	0.95	0.39		
Correlation	0.71			

Table 1B
UK VAR Estimation Results

Dependent Variable	Real T-bill rate	Nominal T-bill rate	Inflation	R ²
VAR Estimation Results				
Real T-bill return	0.07 (0.11)	-0.06 (0.03)	-0.41 (0.10)	0.26
Nominal T-bill yield	0.51 (0.20)	1.055 (0.05)	0.78 (0.19)	0.93
Inflation	0.00 (0.07)	-0.03 (0.02)	0.67 (0.07)	0.87
Const.	-0.0006 (0.0019)	0.0003 (0.0005)	0.0012 (0.0018)	
Observed and hypothetical moments of 10 year real bonds				
	Observed	Hypothetical		
Mean	2.64	2.49		
Standard Deviation	1.00	0.61		
Correlation	0.77			

Table 2
Parameter Estimates for Alternative Risk Models

Parameter	Constant Covariance	Full Model	Persistent Risk
ϕ_x	0.93	0.94	0.95
μ_x	0.0104	0.0028	0.0034
ϕ_v	N/A	0.77	set to 1
μ_v	N/A	-2.01×10^{-5}	0.0010
σ_m	set to 0.23	set to 0.23	set to 0.23
σ_x	0.0031	set to 1	set to 1
σ_{mx}	7.23×10^{-4}	0.23	0.23
$\sigma_{x'}$	N/A	0.0048	0.0031
σ_v	N/A	0.0003	0.0004
β_{em}	N/A	set to 0.41	set to 0.41
σ_{yield}	1.12×10^{-4}	1.16×10^{-6}	9.14×10^{-6}
σ_{cov}	N/A	4.74×10^{-4}	5×10^{-4}
<i>premium</i>	0.0016	0.0157	0.00160

US Inflation Indexed Bonds Outstanding




Figure 1A

UK Inflation Indexed Bonds Outstanding




Figure 1B

US Real and Nominal Yields




Figure 2A

UK Real and Nominal Yields




Figure 2B

US Realized and Breakeven Inflation Rates




Figure 3A

UK Realized and Breakeven Inflation Rates




Figure 3B

Standard Deviations of US Daily Bond Returns Over 1-Year Moving Window




Figure 4A

Standard Deviations of UK Daily Bond Returns Over 1-Year Moving Window

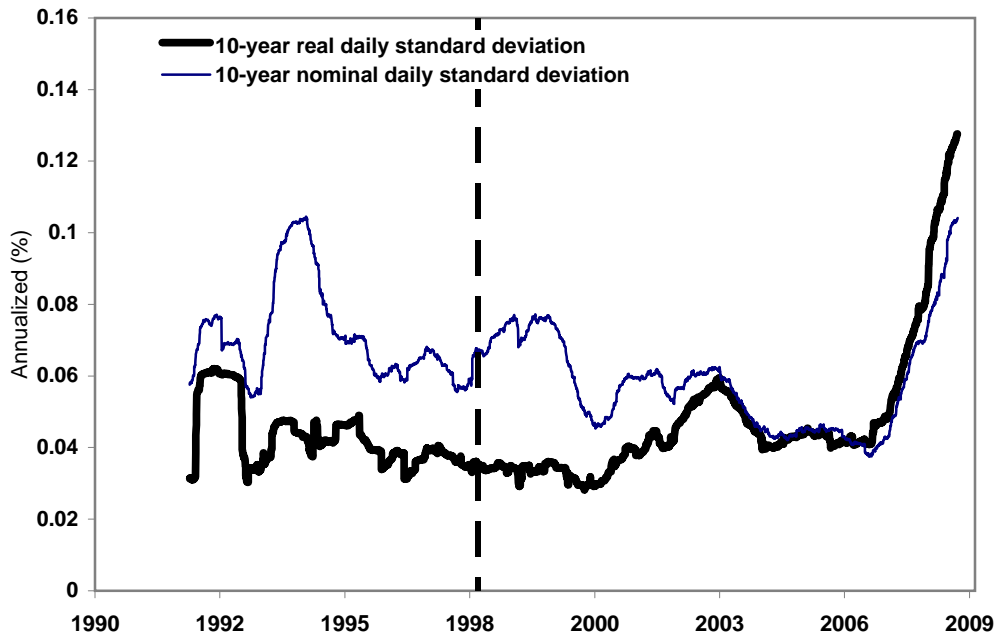


Figure 4B

US Breakeven Inflation Volatility and Nominal/Real Correlation




Figure 5A

UK Breakeven Inflation Volatility and Nominal/Real Correlation




Figure 5B

US Correlations of Daily Bond Returns with Equity Returns




Figure 6A

UK Correlations of Daily Bond Returns with Equity Returns




Figure 6B

US Betas of Daily Bond Returns with Equity Returns




Figure 7A

UK Betas of Daily Bond Returns with Equity Returns

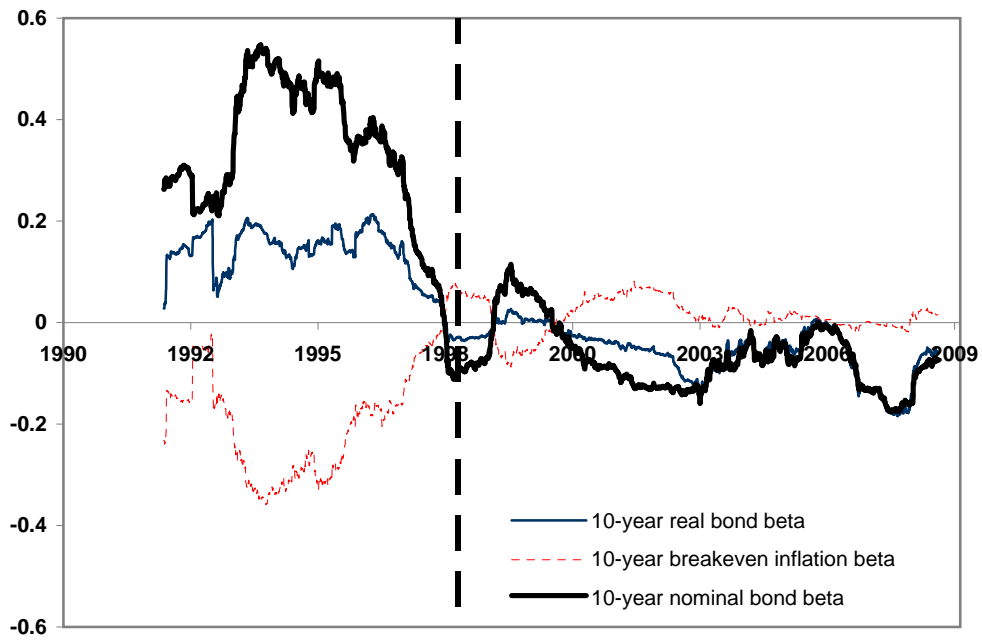


Figure 7B

US Observed and Hypothetical Real Yields




Figure 8A

UK Observed and Hypothetical Real Yields




Figure8B

US 10-Year Real Yields Implied by Alternative Risk Models




Figure 9

US Breakeven Inflation Rates and TIPS Asset Swap Spread




Figure 10




Figure 11




Figure 12