

# Options-Implied Probability Density Functions for Real Interest Rates\*

Jonathan H. Wright  
Johns Hopkins University

This paper constructs options-implied probability density functions for real interest rates. These use options on TIPS, which were launched in 2009. Data availability limits us to studying short-maturity probability density functions for intermediate- to long-term TIPS yields. The PDFs imply high uncertainty about real rates. I also estimate empirical pricing kernels using these option prices along with time-series models fitted to real interest rates. The empirical pricing kernel implies that investors have high marginal utility in states of the world with high real rates.

JEL Codes: C58, E43, G12.

## 1. Introduction

Options markets provide risk-neutral distributions of the future values of many asset prices and even macroeconomic variables. These include probability density functions (PDFs) for future nominal interest rates (see a long literature, including Li and Zhao 2009) and inflation (Kitsul and Wright 2013). But to date, as far as I know, no work has been done on options-implied PDFs for real interest rates. This is the task undertaken in the present paper. A PDF for real interest rates cannot be backed out from PDFs from nominal interest rates and inflation without making an assumption about the correlation between these two. However, over the past seven years, the Chicago Board Options Exchange (CBOE) has traded options on a Treasury Inflation-Protected Securities (TIPS) exchange-traded

---

\*I am grateful to Eric Swanson and John Williams for very helpful comments on an earlier draft. All errors and omissions are my sole responsibility. Author contact: Department of Economics, Johns Hopkins University, Baltimore, MD 21218; wrightj@jhu.edu.

fund (ETF). I show that these can be used to construct risk-neutral PDFs of intermediate- to long-term TIPS yields at short horizons. The standard deviation of options-implied PDFs for nominal interest rates is often treated as a measure of monetary policy uncertainty. Yet, aggregate demand should be determined by real rather than nominal rates, and so PDFs for real rates arguably give a better measure of investors' beliefs about the future effective stance of monetary policy. Even the limited snapshot of uncertainty about real interest rates that can be obtained from the available data may thus be important to monetary policymakers.

The plan for the remainder of this paper is as follows. In section 2, I describe the data and the methodology for extracting risk-neutral PDFs (also known as state-price densities) for real interest rates. Section 3 contains the results for implied volatilities and risk-neutral PDFs. Section 4 evaluates these as physical density forecasts and finds some evidence for the existence of risk premia, notwithstanding the short available data span. Section 5 provides econometric-based physical density forecasts for real interest rates and then adapts the methodology of Rosenberg and Engle (2002) to obtain a time-varying empirical pricing kernel. Based on these results, investors generally appear to view high real interest rate states as having high marginal utility. Section 6 concludes.

## 2. Data and Methodology

Since January 2009, the CBOE has traded call and put options on the Barclays TIPS ETF. These options have expiration dates of up to nine months ahead and various strike prices. The underlying asset is a portfolio of TIPS with a duration that is stable at seven years. I obtained prices (both trades and quotes) from the CBOE from inception to the end of January 2016. Table 1 lists some summary statistics of these options. On average, forty-nine options change hands a day,<sup>1</sup> where each option is for the purchase/sale of 100 shares in the ETF, each of which has a value of roughly \$100. Thus the daily transaction volume has a notional underlying amount of about \$500,000. This amount is minuscule relative to trading in

---

<sup>1</sup>This is the total number of options that change hands each day—the number of distinct trades is smaller.

**Table 1. Options Summary Statistics**

	Mean	SD	Min.	Max.
<i>Trades</i>				
Daily Call Volume	28	106	0	2,528
Daily Put Volume	21	71	0	1,175
<i>Quotes</i>				
Moneyness (Call)	0.94	0.07	0.57	1.14
Moneyness (Put)	1.03	0.06	0.79	1.34
Maturity	102	73	1	243
Bid-Ask Spread (%)	18.4	11.2	0.9	50.0
<p><b>Notes:</b> Volume represents the total volume per day in number of contracts. Each contract is for 100 units in the ETF. Moneyness is the ratio of the strike price to the underlying price for all trades. Maturity is measured in days. Bid-ask spread is measured as the ask price less the bid price, divided by the midpoint, and multiplied by 100 to be in percentage points.</p>				

conventional interest rate derivatives—for comparison, trading volume in Eurodollar options averaged around 440,000 contracts per day in 2015. But that does not mean that the prices are uninformative. In experimental game theory, it is common to study individuals' behavior when tiny sums are at stake. Prediction markets are a good example of cases in which the stakes are small, but there is strong evidence that prices are nonetheless quite informative (Wolfers and Zitzewitz 2004). The dollar trading volumes in TIPS options are far bigger than those in experimental games and in prediction markets, and it seems reasonable to presume that they are reflective of traders' beliefs. Nonetheless, it is clear that this is a small and illiquid market and results should be treated with caution, especially in 2009 and 2010, when liquidity was particularly limited.

I do all empirical work with midpoints of bid and ask quotes. This makes for a large data set of  $n = 164,320$  options quotes. I prefer this to using actual trades because that would give a much smaller data set and one in which bid-ask bounce would be a major concern. Bid-ask spreads are wide and average around 18 percent. The quotes on call/put options have average moneyness slightly below/above one, respectively, indicating that the options quotes tend to be in the money. Note, however, that the price of an in-the-money put/call

option can be converted to that of an out-of-the-money call/put option via put-call parity, which is an arbitrage relationship (Stoll 1969).

### 2.1 *Semiparametric State-Price Densities*

In estimating state-price densities, one possible approach is to use the observed quotes on options at a particular maturity on a particular day in isolation to directly estimate the state-price density at that maturity on that day, as in Bliss and Panigirtzoglou (2002) and Kitsul and Wright (2013). In the current context, I do not use that approach for a number of reasons. Firstly, there are on average twenty-one strike prices for each maturity/trading day combination, which constitutes a very small sample size, especially given that the data are noisy. Secondly, there are gaps in the available strike prices in many cases. Thirdly, the available options are for fixed maturity dates (the third Friday of the month), whereas it is easier to compare state-price densities over time when they are for fixed horizons, like three months ahead.

Instead, I use a smoothing approach to estimate prices that are not actually observed. Intuitively, the idea is to infer prices from the observed prices of options that are similar in maturity, strike price, and other characteristics. This smooths out measurement error in individual quotes. More precisely, I adopt the semiparametric approach proposed by Ait-Sahalia and Lo (1998). I convert the observed options prices into Black-Scholes implied volatilities. Let these options-implied volatilities be  $\{\sigma_i\}_{i=1}^n$ , where  $n$  is the total number of options, pooling all call and put options on all trading days, strike prices, and expiration dates. I assume that the implied volatility is a function of a state vector,  $\sigma(Z_i)$ , where the  $p \times 1$  state vector,  $Z_i$ , includes the strike price or moneyness of the option. I assume that the price of a call option is  $C_{BS}(S, X, r_f, d, \tau, \sigma(Z))$ , where  $C_{BS}(\cdot)$  is the Black-Scholes expression for the price of a call option,  $S$  is the price of the underlying security,  $X$  is the strike price of the option,  $r$  is the risk-free rate,  $d$  is the dividend yield of the security, and  $\tau$  is the time to maturity of the option. I estimate the implied volatility function nonparametrically as

$$\hat{\sigma}(Z) = \frac{\sum_{i=1}^n \prod_{j=1}^p k\left(\frac{Z_j - Z_{ij}}{h_j}\right) \sigma_i}{\sum_{i=1}^n \prod_{j=1}^p k\left(\frac{Z_j - Z_{ij}}{h_j}\right)}, \quad (1)$$

where  $k(\cdot)$  is the Gaussian kernel and the bandwidth  $h_j$  is set to  $\frac{c_0 \sigma_j n^{-1/11}}{\log(n)}$ , where  $\sigma_j$  is the standard deviation of  $Z_{ij}$ . The value of  $c_0$  is determined by tenfold cross-validation.<sup>2</sup> The nonparametric estimate of the implied volatility gives an implied call option price for any choice of  $Z, S, r_f, d$ , and  $\tau$ . The second derivative of this call price is then the state-price density, or risk-neutral options-implied PDF (Breedon and Litzenberger 1978). In the empirical work, the risk-free rate is measured from three-month LIBOR—over the period in question, short-term risk-free rates are very close to zero. Meanwhile, the dividend yield is measured as the trailing twelve-month yield on the TIPS ETF.

Equation (1) can be estimated over the full sample, effectively pooling information across the whole sample period. Alternatively, equation (1) can be estimated for each calendar month separately. This pools information across multiple days but allows for the possibility of lower frequency changes that are not captured by the state vector. I have computed the PDFs in both ways. However, in order to conserve space, I report the result only using the approach where the implied volatility is estimated for each calendar month separately.

This of course gives us a PDF for the price of a share in the TIPS ETF. But this ETF has a duration of about seven years, giving the approximation

$$\Delta y_t \simeq -\frac{1}{7} \Delta \log(P_t), \quad (2)$$

where  $P_t$  is the price of the ETF and  $y_t$  is the zero-coupon TIPS yield. This in turn allows the PDF for the TIPS price to be converted into a PDF for the TIPS yield. I use the data set of Gürkaynak, Sack, and Wright (2010) for zero-coupon TIPS yields. It should be noted that equation (2) is empirically a very good approximation—regressing three-month changes in seven-year zero-coupon yields on

---

<sup>2</sup>Standard, or leave-one-out, cross-validation would pick  $c_0$  so as to minimize the mean square fitting error, where the fit for each observation is computed omitting that one observation. This would be very computationally costly given the large sample size that I have. Instead, I split the data into ten roughly equal-sized bins, and compute the errors within each bin, where the fit is computed omitting that bin. The parameter  $c_0$  is chosen to minimize the mean square fitting error over all these bins. This reduces the number of nonparametric function estimates from the sample size to ten.

the corresponding change in the log price of the TIPS ETF gives a coefficient estimate of almost exactly  $-\frac{1}{7}$  ( $-0.159$ ) with an R-squared of 93 percent.


Moreover, as the vast majority of changes in TIPS yields are level shifts, equation (2) can also provide a good approximation to the PDF for the change in any TIPS yield with maturity in at least the five- to ten-year range. The information that can be gathered from these TIPS options is limited by the fact that they are short-maturity options on longer-term TIPS yields, but this nonetheless represents an important first look at state-price densities for real interest rates.

### 3. Results


#### 3.1 *Implied Volatilities*

I first plot the average annualized implied volatility of TIPS prices, by month, averaged across all options with moneyness between 0.9 and 1.1, in figure 1. The implied volatility was high in early 2009. This was in the immediate aftermath of the financial crisis—TIPS prices were volatile and liquidity was unusually poor. Implied volatilities climbed in 2011, which was during the European debt crisis. Since 2012, TIPS implied volatilities have been fairly stable. Over the sample period, TIPS implied volatilities have averaged around 10 percentage points at an annualized rate. Given the relationship between real yields and prices (equation (2)), this means that the options-implied volatility of real rates is around 1.4 percentage points per annum. Using daily real yield data (Gürkaynak, Sack, and Wright 2010), the realized volatility of seven-year TIPS zero-coupon yields over the sample period is 80 basis points per annum. In other financial markets, it is standard to find that options-implied volatilities exceed realized volatilities, and this is interpreted as a variance risk premium (Bollerslev, Tauchen, and Zhou 2009; Choi, Mueller, and Vedolin 2013). I observe the same phenomenon for TIPS yields.

To compare these implied volatilities with their nominal counterparts, I used the same methodology to obtain the annualized implied volatility of options on ten-year nominal Treasury futures—a long-standing and very liquid options market. These are shown in figure 2. The options-implied volatility of a ten-year nominal futures

**Figure 1. Implied Volatility of TIPS Options**

**Notes:** This figure plots the annualized volatility of TIPS implied by the Black-Scholes formula averaged over all options within a month, using only options with moneyness between 0.9 and 1.1.

**Figure 2. Implied Volatility of Nominal Treasury Options**

**Notes:** This figure plots the annualized volatility of ten-year nominal Treasury futures contracts implied by the Black-Scholes formula averaged over all options within a month, using only options with moneyness between 0.9 and 1.1.

contract was also elevated in early 2009, but not as much as the TIPS implied volatility. This is perhaps because the financial crisis led to fewer strains in the market for nominal Treasury securities. The options-implied volatility of a ten-year nominal futures contract fell in the first part of 2011, while TIPS implied volatilities were rising, though the nominal implied volatility climbed a bit later that year. The effect of the “taper tantrum” of mid-2013 is more apparent in nominal implied volatilities than in TIPS implied volatilities. That’s perhaps not surprising, since Federal Reserve large-scale asset purchases were overwhelmingly concentrated in nominal securities. Nominal implied volatilities also climbed in late 2014, as oil prices slid, and this did not affect TIPS implied volatilities very much.

Over the sample period, nominal implied volatilities have averaged around 7 percentage points at an annualized rate. As discussed in Burghardt and Belton (2005), the precise terms of a ten-year nominal Treasury futures contract are such that the effective underlying security is a nominal coupon security with 6.5 to 7 years to maturity (the “10-year” label is somewhat misleading). The duration of this security is around six years. So the options-implied volatility of six-year zero-coupon nominal yields is around 1.2 percentage points per annum. Meanwhile, using daily nominal yield data, the realized volatility of six-year nominal zero-coupon yields over the sample period is 92 basis points per annum.


As Campbell, Shiller, and Viceira (2009) point out, whereas one might in theory think that TIPS yields ought to be determined by the long-run marginal product of capital and ought to be very stable, in practice TIPS yields are quite unstable. This shows up in the options market as well. The spread between options-implied and realized volatility (the variance risk premium) seems to be bigger for TIPS than for nominal bonds.

Figure 3 plots the average TIPS implied volatility by moneyness of the option. Under the Black-Scholes assumptions, this should be flat. But it in fact exhibits a pronounced volatility smile.

### *3.2 Computing PDFs*

I next turn to the construction of the full PDFs. I consider two choices of the state vector that governs the implied volatility of different options, as follows:



**Figure 3. Implied Volatility of TIPS Options**

**Notes:** This figure plots the annualized volatility of TIPS implied by the Black-Scholes formula averaged across the whole sample by moneyness bin. The moneyness is the ratio of the strike price to the underlying price.

- S1: The state vector is  $(X, S, \tau)'$ —the strike price, price of the underlying security and maturity.
- S2: The state vector is  $(X/S, \tau, \sigma_{EGARCH})'$ —the moneyness and maturity of the option, and the volatility of daily changes in seven-year TIPS yields implied by an EGARCH(1,1) model with conditionally  $t$ -distributed errors. The parameters of the EGARCH model are estimated by maximum likelihood over the whole period for which TIPS yields are available (since January 1999).

The idea of using EGARCH-filtered volatility as an element of the state vector was used in the context of nominal interest rate derivatives by Li and Zhao (2009).

Table 2 shows the options pricing errors using either of these two state vectors—the percentage difference between actual options prices and the fitted options price using the nonparametric estimate of the volatility function in equation (1) (estimated for each month separately, as discussed above). The median absolute pricing error is a bit above 2 percent with either state vector.  $R^2$  values—the

**Table 2. Options Fitting Errors**

State Variables	S1	S2
Median Absolute Error	2.38	2.06
R-squared	75.6	82.0

**Notes:** The median absolute error is the median absolute difference between actual and fitted options prices, in percentage points. The fitted options prices are Black-Scholes prices, with volatility given by equation (1). The R-squared is the ratio of the variance of fitted implied volatilities to observed implied volatilities, in percentage points.

ratio of the variance of fitted implied volatilities to observed implied volatilities—are around 80 percent. The methodology is using fitted options prices, so it is an important check on the adequacy of the methodology and state vector that the options pricing errors not be too big. It can be seen from table 2 that this is indeed the case.

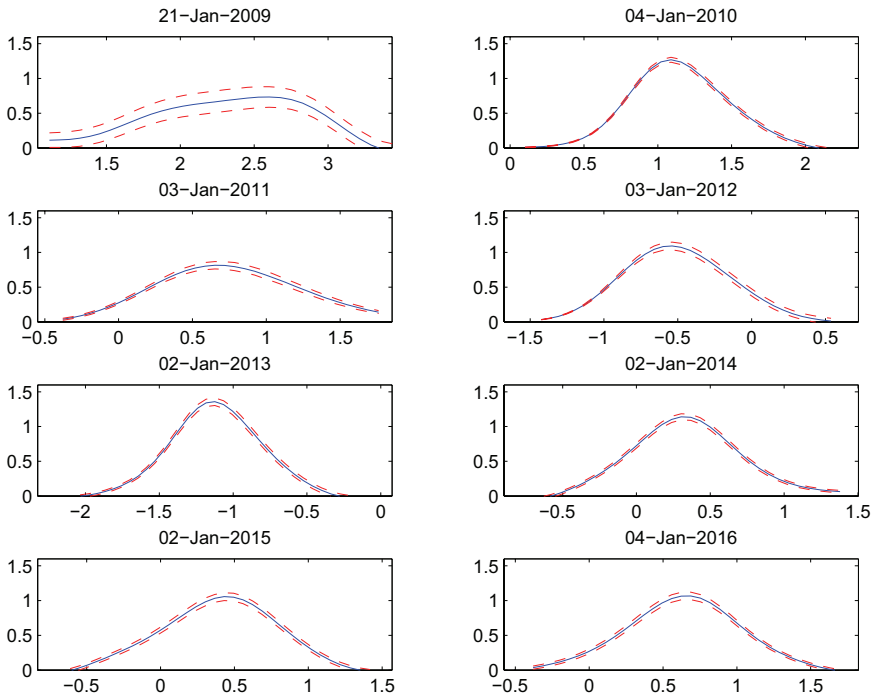
State variables S1 and S2 both contain three elements and are thus equally parsimonious. State vector S1 can be thought of as pooling options with similar maturity and moneyness over days when the level of real rates is similar. State vector S2 can instead be thought of as pooling options with similar maturity and moneyness over days when the level of real rate volatility is similar. Both of these are reasonable modeling choices. From table 2 we see that in the sample at hand, both state vectors provide a good fit to observed options prices, although S2 does a bit better.

Figures 4 and 5 show the implied PDFs for real yields at the three-month maturity as of the first day of each year from 2009 to 2016, inclusive, for the two different state vectors. The densities have a standard deviation of real rates that is around 60 basis points. The densities are more dispersed in the early years of the sample (2009–11) than more recently, especially when using state vector S2. Asymptotic 95 percent confidence intervals are included in these figures.<sup>3</sup> The uncertainty in the PDFs is driven by measurement error in options prices. While measurement error for individual

---

<sup>3</sup>These are formed from the variance-covariance matrix of implied volatilities, computed assuming i.i.d. measurement error in options prices and then applying the delta method.

**Figure 4. Three-Month-Ahead PDF for Real Rates:  
State Variable S1**



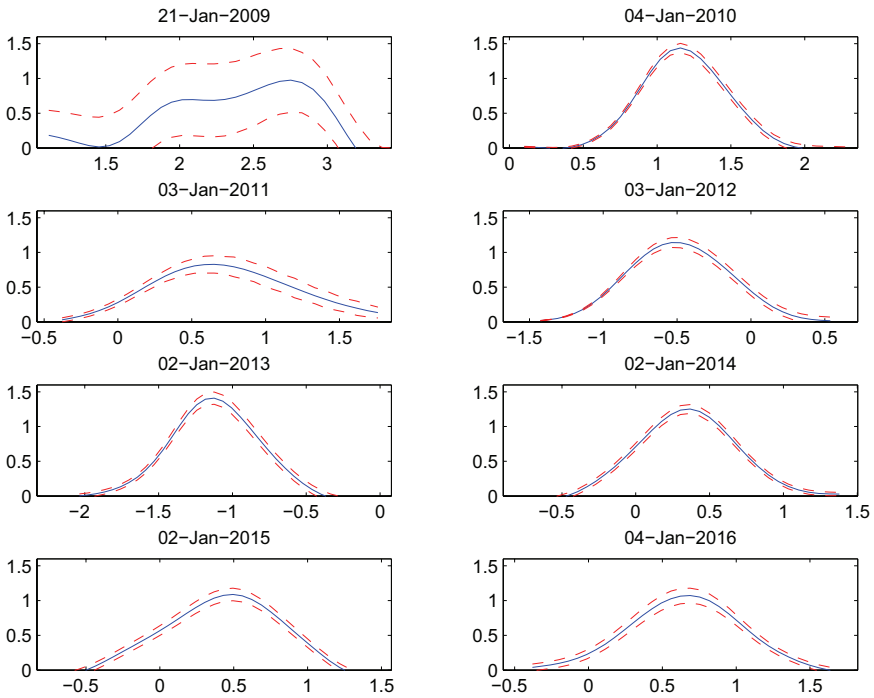
**Notes:** This figure plots the three-month-ahead PDF for seven-year TIPS yields constructed by the method described in the paper using the state variable S1. The PDFs are shown as of the first day in each year. Dashed lines are asymptotic 95 percent confidence intervals. The estimation of the implied volatility function in equation (1) is done for each calendar month separately.

options prices is considerable, with the smoothing procedure it is overwhelmed by the large sample size, leaving tight confidence intervals. The prices used to construct the state-price density are these smoothed prices that average out measurement error in individual quotes.<sup>4</sup>

Figures 6 and 7 repeat the exercise at the six-month maturity. Again, the densities are more dispersed early in the sample,

<sup>4</sup>If equation (1) were estimated over the entire sample, the standard errors would be even tighter.

**Figure 5. Three-Month-Ahead PDF for Real Rates:  
State Variable S2**



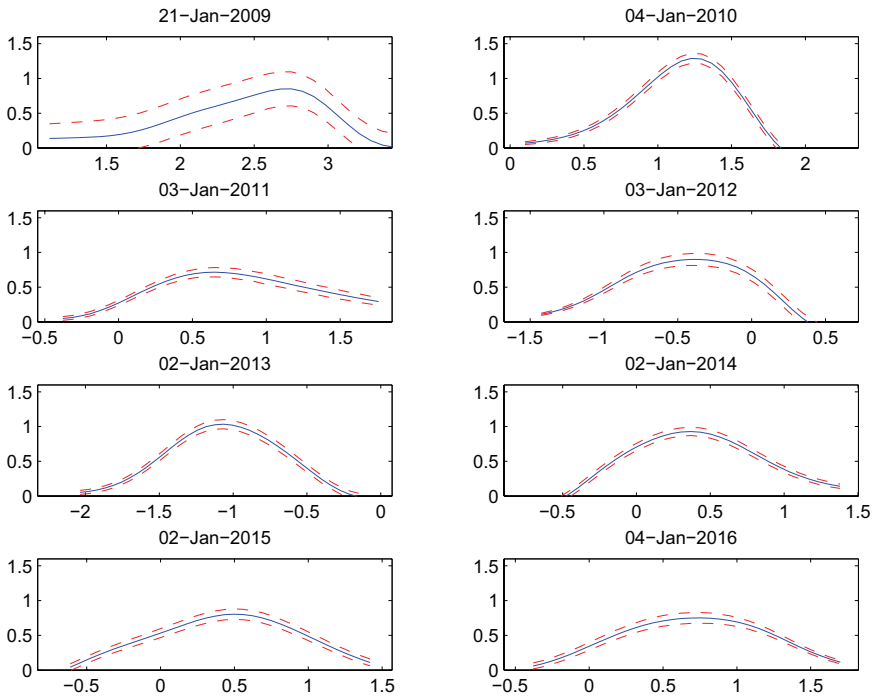
**Notes:** This figure plots the three-month-ahead PDF for seven-year TIPS yields constructed by the method described in the paper using the state variable S2. The PDFs are shown as of the first day in each year. Dashed lines are asymptotic 95 percent confidence intervals. The estimation of the implied volatility function in equation (1) is done for each calendar month separately.

especially when using state vector S2, and are precisely estimated. The densities are modestly skewed in the direction of higher real rates.

#### 4. Evaluation of Density Forecasts

If agents were risk neutral, then the PDFs reverse-engineered from options ought to be good density forecasts for future TIPS yields. As the market in TIPS options is still young, evaluation of these options-implied density forecasts is difficult. Still, I have

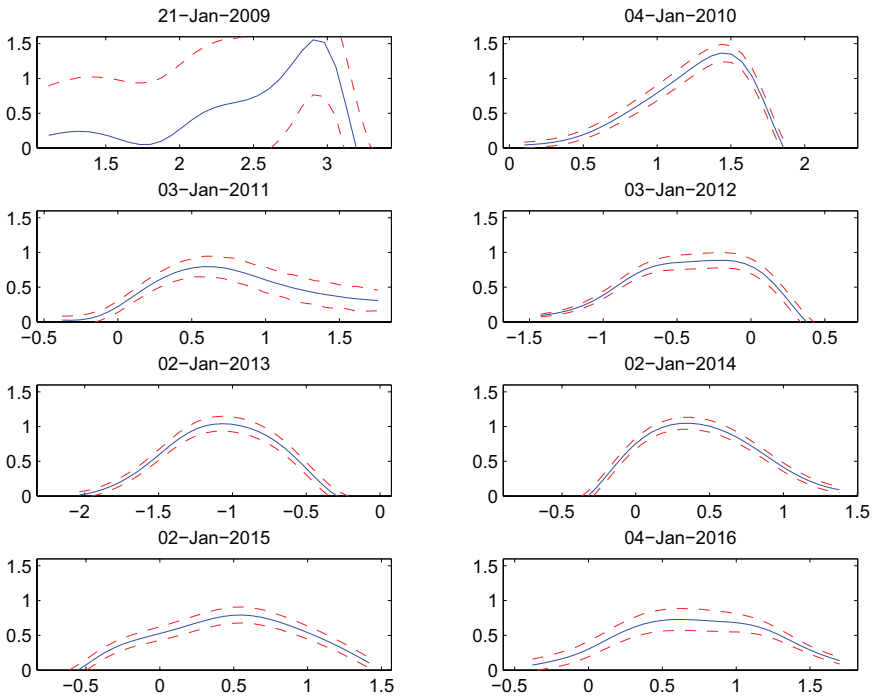
**Figure 6. Six-Month-Ahead PDF for Real Rates:  
State Variable S1**



**Notes:** This figure plots the six-month-ahead PDF for seven-year TIPS yields constructed by the method described in the paper using the state variable S1. The PDFs are shown as of the first day in each year. Dashed lines are asymptotic 95 percent confidence intervals. The estimation of the implied volatility function in equation (1) is done for each calendar month separately.

twenty-eight non-overlapping one-quarter-ahead density forecasts, which is enough to make an attempt at density forecast evaluation. I follow Diebold, Gunther, and Tay (1998) in evaluating these density forecasts by checking that the probability integral transform (PIT) of the realized data is both uniform on the unit interval and independent over time. Figure 8 shows the histogram of the PIT of the realized data for non-overlapping forecasts of one-quarter-ahead TIPS yields made at the end of each quarter from 2009:Q1 to 2015:Q4, respectively, using state variable S1. The realized real rate tends to


**Figure 7. Six-Month-Ahead PDF for Real Rates:  
State Variable S2**



**Notes:** This figure plots the six-month-ahead PDF for seven-year TIPS yields constructed by the method described in the paper using the state variable S2. The PDFs are shown as of the first day in each year. Dashed lines are asymptotic 95 percent confidence intervals. The estimation of the implied volatility function in equation (1) is done for each calendar month separately.

lie toward the left of the forecast density, consistent with a positive real term premium. The likelihood-ratio test of Berkowitz (2001)—which jointly tests that the PIT is both uniform and independent over time—rejects with a  $p$ -value of below 0.01. Thus, despite a small sample, I can reject the hypothesis that the options-implied PDFs embeds no risk premium at conventional significance levels.

The finding that the options-implied PDFs embed a risk premium, even though statistically significant, could nevertheless be idiosyncratic to this unusual period. Swanson and Williams (2014)

**Figure 8. Histogram of the Probability Integral Transform**

**Notes:** This figure shows the histogram of the probability integral transform of three-month-ahead density forecasts of TIPS yields. As described in Diebold, Gunther, and Tay (1998), with an optimal density forecast, these should be uniform on the unit interval in population. The density forecasts are made at the end of each quarter from 2009:Q1 to 2015:Q4, for a total of twenty-eight forecasts, using state variable S1.

document that interest rate futures and surveys both consistently predicted near-term liftoff from the zero lower bound throughout much of the period from 2009 to 2015, even though that didn't actually happen until the end of 2015. Interest rate futures embed a term premium, but surveys should not. Time alone will tell whether realized real rates continue to lie toward the left of the forecast density or whether this was idiosyncratic to the years following the Great Recession.

## 5. Physical PDFs and the Empirical Pricing Kernel

Probability densities under the physical ( $P$ ) measure weight different outcomes by their actual probabilities. The implied PDFs extracted from options reweight outcomes by marginal utility, and so give probabilities under the risk-neutral ( $Q$ ) measure. Outcomes

in which the marginal utility of an investor is high get more weight under the risk-neutral measure than under the physical measure. The evidence from the last section indicates that physical and risk-neutral PDFs are significantly different, even with the short sample period that is available.

In this section, I estimate physical PDFs for real interest rates and adapt the methodology of Rosenberg and Engle (2002) to obtain a time-varying empirical pricing kernel projected onto real interest rate states.

First, I fit a univariate EGARCH(1,1) model with  $t$ -distributed conditional errors to daily changes in seven-year TIPS yields. I estimate the model over the whole period for which TIPS yields are available by maximum likelihood, exactly as in the construction of state variable S2 above. Then, at each date, I simulate density forecasts for these TIPS yields. If the model is correctly specified, then these are density forecasts under the physical measure. Figure 9 shows these EGARCH-based density forecasts at the start of each year at the three-month horizon; results at the six-month horizon are not shown, but are qualitatively similar. This specification does not allow for mean reversion in real interest rates, but it does allow for potentially asymmetric conditional volatility and conditional excess kurtosis. These latter features are most important for near-term density forecasts.

Let  $r_t$  denote the seven-year real interest rate. I specify a pricing kernel from time  $t$  to time  $t + \tau$  that is a nonlinear function of  $r_{t+\tau}$ :

$$M_t(\theta_t, r_{t+\tau}) = \theta_{0t} \exp(\theta_{1t}T_1(r_{t+\tau}) + \theta_{2t}T_2(r_{t+\tau})), \quad (3)$$

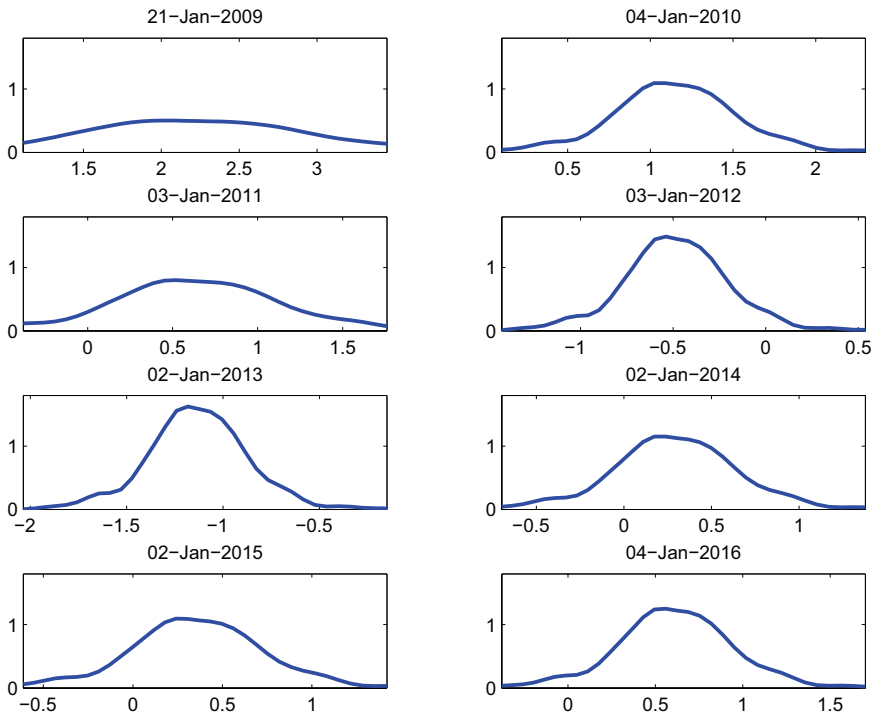
where  $\theta_t = (\theta_{1t}, \theta_{2t})'$  is a vector of parameters,  $\theta_{0t} = \frac{e^{-r_{ft}(\tau)\tau}}{E(\exp(\theta_{1t}T_1(r_{t+\tau}) + \theta_{2t}T_2(r_{t+\tau})))}$ ,  $r_{ft}(\tau)$  is the nominal  $\tau$ -period risk-free rate, and  $T_j(\cdot)$  is a Chebyshev polynomial, such that  $T_j(x) = \cos(j \cos^{-1}(\frac{2x-a-b}{b-a}))$ , where  $a$  and  $b$  are the endpoints of the interval over which the Chebyshev polynomial is expected to provide an approximation.<sup>5</sup> These are set to 1.2 percentage points below and above the current real interest rate,  $r_t$ . At each time  $t$ , this gives an empirical pricing kernel.

---

<sup>5</sup>The specification of  $\theta_{0t}$  ensures that the expectation of the pricing kernel is the inverse of the price of a  $\tau$ -period risk-free bond.



**Figure 9. Physical PDF for Real Rates:  
Three-Month Horizon**



**Notes:** This figure plots the three-month-ahead PDF for seven-year TIPS yields constructed from fitting a  $t$ -EGARCH model to daily changes in seven-year TIPS yields, as described in the text. The PDFs are shown as of the first day in each year.

Let  $f_t^Q(r_{t+\tau})$  denote the risk-neutral PDF for  $r_{t+\tau}$  as of time  $t$ , obtained in section 3. The horizon  $\tau$  is either three or six months, and the state vector can be either S1 or S2. Let  $f_t^P(r_{t+\tau})$  denote the physical PDF obtained from the fitted EGARCH model. As is well known,

$$f^Q(r_{t+\tau}) = e^{r_{f,t}(\tau)} f^P(r_{t+\tau}) M_t, \tag{4}$$

where  $M_t$  is the pricing kernel from time  $t$  to time  $t + \tau$ . I estimate the parameters of the pricing kernel to solve the optimization problem:

$$\hat{\theta}_t = \arg \min_{\theta_t} \sum_{i=1}^L (f^Q(r_{i,t+\tau}) - e^{r_{f,t}(\tau)} f^P(r_{i,t+\tau}) M_t(\theta_t, r_{i,t+\tau}))^2, \quad (5)$$

where  $\{r_{i,t+\tau}\}_{i=1}^L$  is a grid of possible values of  $r_{t+\tau}$ . The grid consists of 32 points in the interval  $r_t \pm 1.2$  percentage points.

Intuitively, this procedure is simply measuring the empirical pricing kernel as the ratio of the  $Q$ -measure PDF to the  $P$ -measure PDF. However, following the approach of Rosenberg and Engle (2002), I restrict the pricing kernel to have the functional form in equation (3).

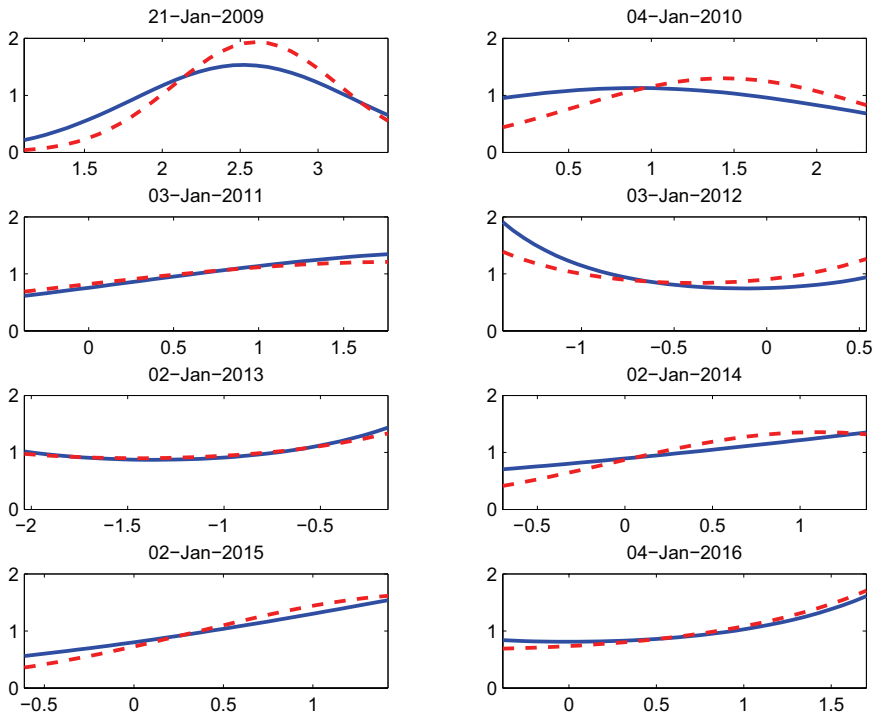
Figure 10 shows the estimated empirical pricing kernels at the start of each year, at the horizon,  $\tau$ , of three months. Although the pricing kernels vary over time, the empirical pricing kernels generally slope up, meaning that marginal utility is highest in high real interest rate states of the world. As in the previous section, this is consistent with a positive real term premium. It contrasts somewhat with estimated pricing kernels for stock returns or inflation, which are found to be U-shaped in returns and inflation (see, for example Ait-Sahalia and Lo 2000, Jackwerth 2000, and Kitsul and Wright 2013).

## 6. Conclusion

Options on TIPS have recently been launched, and in this paper I have used these to construct risk-neutral probability densities for real interest rates going back to 2009. The nature of the options that trade means that it is only possible to obtain PDFs at horizons of a few months and for real interest rates at the seven-year maturity. This contrasts with the situation for nominal interest rates, where there is a wide range of expiration times and underlying maturities.

The PDFs that I obtained imply high uncertainty about real rates, especially in 2009–11. As is well known, options-implied PDFs embed risk premia, and I have shown direct evidence that they should not be interpreted as physical density forecasts. Risk-neutral PDFs overweight high marginal utility states of the world, relative to their actual physical probabilities. The high marginal utility states of the world appear to be associated with high real interest rates. However, Feldman et al. (2015) make a strong case that policymakers

**Figure 10. Empirical Pricing Kernels:  
Three-Month Horizon**



**Notes:** This figure plots the three-month empirical pricing kernels as a function of seven-year TIPS yields, constructed as described in section 5. Results are shown using state variable S1 (solid line) and state variable S2 (dashed line). The empirical pricing kernels are shown as of the first day in each year.

should actually *want* risk-neutral probabilities, weighted by marginal utility, rather than physical probabilities.

I hope that the PDFs for index-linked bond yields provided here can be useful to central banking practice, and can be used and analyzed more extensively as the market in TIPS options matures. Since 2009, TIPS and nominal interest rate implied volatilities have diverged at times, and they might do so more in the future. Uncertainty about future real interest rates can now—subject to many limitations—be directly reverse-engineered from asset prices.

## References

- Aït-Sahalia, Y., and A. W. Lo. 1998. "Nonparameteric Estimation of State-Price Densities Implicit in Financial Asset Prices." *Journal of Finance* 53 (2): 499–547.
- . 2000. "Nonparametric Risk Management and Implied Risk Aversion." *Journal of Econometrics* 94 (1–2): 9–51.
- Berkowitz, J. 2001. "Testing Density Forecasts, with Applications to Risk Management." *Journal of Business and Economic Statistics* 19 (4): 465–74.
- Bliss, R., and N. Panigirtzoglou. 2002. "Testing the Stability of Implied Probability Density Functions." *Journal of Banking and Finance* 26 (2–3): 381–422.
- Bollerslev, T., G. Tauchen, and H. Zhou. 2009. "Expected Stock Returns and Variance Risk Premia." *Review of Financial Studies* 22 (11): 4463–92.
- Breeden, D. T., and R. H. Litzenberger. 1978. "Prices of State Contingent Claims Implicit in Option Prices." *Journal of Business* 51 (4): 621–51.
- Burghardt, G., and T. Belton. 2005. *The Treasury Bond Basis: An In-Depth Analysis for Hedgers, Speculators and Arbitrageurs*. McGraw-Hill.
- Campbell, J. Y., R. J. Shiller, and L. Viceira. 2009. "Understanding Inflation-Indexed Bond Markets." *Brookings Papers on Economic Activity* (Spring): 79–120.
- Choi, H., P. Mueller, and A. Vedolin. 2013. "Bond Variance Risk Premiums." Working Paper.
- Diebold, F. X., T. A. Gunther, and A. S. Tay. 1998. "Evaluating Density Forecasts with Applications to Financial Risk Management." *International Economic Review* 39 (4): 863–83.
- Feldman, R., K. Heinecke, N. Kocherlakota, S. Schulhofer-Wohl, and T. Tallarini. 2015. "Market-Based Probabilities: A Tool for Policymakers." Working Paper, Federal Reserve Bank of Minneapolis.
- Gürkaynak, R. S., B. Sack, and J. H. Wright. 2010. "The TIPS Yield Curve and Inflation Compensation." *American Economic Journal: Macroeconomics* 2 (1): 70–92.
- Jackwerth, J. C. 2000. "Recovering Risk Aversion from Option Prices and Realized Returns." *Review of Financial Studies* 13 (2): 433–51.

- Kitsul, Y., and J. H. Wright. 2013. "The Economics of Options-Implied Inflation Probability Density Functions." *Journal of Financial Economics* 110 (3): 696–711.
- Li, H., and F. Zhao. 2009. "Nonparametric Estimation of State-Price Densities Implicit in Interest Rate Cap Prices." *Review of Financial Studies* 22 (11): 4335–76.
- Rosenberg, J. V., and R. F. Engle. 2002. "Empirical Pricing Kernels." *Journal of Financial Economics* 64 (3): 341–72.
- Stoll, H. R. 1969. "The Relationship Between Put and Call Option Prices." *Journal of Finance* 24 (5): 801–24.
- Swanson, E. T., and J. Williams. 2014. "Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates." *American Economic Review* 104 (10): 3154–85.
- Wolfers, J., and E. Zitzewitz. 2004. "Prediction Markets." *Journal of Economic Perspectives* 18 (2): 107–26.