Discussion of "Options-Implied Probability Density Functions for Real Interest Rates"

Eric T. Swanson
University of California, Irvine

Options data provide us with substantial information about the future price of an asset. For example, researchers can use interest rate options to infer the market-implied risk-neutral probability density function (PDF) of the underlying interest rate at expiration (see, e.g., Bank for International Settlements 1999; Bliss and Panigirtzoglou 2002). These option-implied PDFs provide useful information about market participants' views of the likely outcomes for the underlying interest rate, and the relative likelihood of those different outcomes.[1]

Options on nominal interest rates have been widely traded for over thirty years. More recently, the introduction of options on inflation swaps has allowed researchers to estimate market-implied risk-neutral PDFs for inflation as well (Kitsul and Wright 2013). However, unlike with futures data, there is no simple way to combine the PDF for inflation with the PDF for the nominal interest rate to obtain the PDF for the real interest rate. If inflation and the nominal interest rate are correlated (as they would be under a Taylor rule, for example), then this correlation will be important for the PDF of the real interest rate, but it cannot be inferred from the individual PDFs for inflation and the nominal interest rate.

In "Options-Implied Probability Density Functions for Real Interest Rates," Jonathan Wright introduces us to newly available data on options on U.S. Treasury Inflation-Protected Securities (TIPS)—in particular, on a basket of TIPS held by an exchange-traded mutual fund—and uses these to compute market-implied

[1]It's important to note that these PDFs are computed under the risk-neutral measure—hence the term "market-implied risk-neutral PDF"—so they do not quite represent actual probabilities. Instead, they are the probabilities that are implied by the observed options prices, assuming that market participants are risk neutral.

risk-neutral PDFs for the real interest rate directly.[2] This marks the first time anyone has computed a market-implied PDF for the U.S. real interest rate. The results are summarized in figure 4 of Jonathan's paper. The implied risk-neutral PDFs have a fairly standard hump shape with fat tails, and a variance that appears to have changed substantially over time.

1. Data Limitations

The primary contribution of Jonathan's paper—the introduction and analysis of a new data set—carries with it some important limitations related to the novelty of that data. First, these TIPS options are still very lightly traded. On an average day in Jonathan's sample, only forty-nine options contracts in total changed hands. (Note that this is the total trading volume; the number of trades is necessarily less than or equal to this.) In comparison, total trading volume in Eurodollar options averaged about 400,000 contracts per day in 2015, and total trading volume in Treasury note options was of a similar magnitude.[3] This implies that the Chicago Board Options Exchange TIPS options data are about 10,000 times more sparse than the interest rate options data that many researchers are more familiar with and may have used in the past. Jonathan correctly points out that the relatively low trading volumes in the new data set do not imply that the prices are uninformative, particularly if we focus on bid and ask quotes rather than actual executed trades. (By focusing on the midpoint of bid and ask quotes, the number of observations in Jonathan's sample increases to about ninety-five per day, on average.) Nevertheless, low trading volumes imply that prices will tend to be noisier and less informative than they would be in a thicker market.

Low trading volumes also tend to produce large bid-ask spreads, and that is the case for these TIPS options as well. As Jonathan notes, bid-ask spreads average about 18 percent in these data.

[2]To be precise, the PDFs are for the implied yield to maturity on the basket of TIPS underlying the mutual fund.

[3]See the CME Group's Monthly Average Daily Volume Report, available at http://www.cmegroup.com/market-data/volume-open-interest.html.

Jonathan's analysis focuses on the midpoints of the bid and ask quotes, but nevertheless the "true" shadow price could lie anywhere in between, in theory. This suggests that there could be measurement error of the true shadow price of each contract on the order of ± 9 percent.

Finally, note that bid and ask quotes are not available for every strike price on every day of the sample. Quotes seem to be particularly sparse earlier in the sample (in 2009 and 2010), and many of the quotes that do exist then are for deep in- or out-of-the-money strikes, which convey very little information. Thus, one should treat the earlier years of Jonathan's sample with particular caution.

2. Standard Errors

Given the significant limitations of the data, discussed above, it's natural to wonder why Jonathan's reported standard errors in his figures 4–7 are so small. After all, each data point in the sample is plausibly contaminated with a measurement error of ± 9 percent, and that error alone would seem to suggest wider standard-error bands than are reported in the figures.

The difference is due to the type of standard errors that Jonathan is reporting. In particular, they are not the standard errors that would result from analyzing a single day of data in isolation. Instead, the standard errors are what one gets from pooling data on the given day with data from all other days that are *similar to* the given day within the month. (I discuss this pooling and Jonathan's non-parametric estimation procedure in more detail below.) Because there are effectively many days that are similar to any given day in Jonathan's data, the sample size is effectively larger—as if we had roughly 2,000 observations instead of 95. Thus, the standard errors are correspondingly smaller, because the effective sample size is larger.

In other words, the standard errors are computed under the assumption that the given day has many other days that are just like it. The reported standard errors are for this large set of days, all of which can be thought of as being essentially identical to the given day.

3. Estimation Methodology for Probability Density Functions

A related question is whether Jonathan's non-parametric estimation strategy of the implied volatility function is the best approach (see equation (1) of his paper). When working with heavily traded options data, such as Eurodollar options, it is standard practice to estimate the market-implied risk-neutral PDF for each day in isolation. In other words, treat each day as independent from the others in the sample, and estimate the market-implied PDF using the cross-section of options prices available on that day.

TIPS options do not have nearly as many price observations as do Eurodollar options; nevertheless, there may be more than enough observations per day to make the cross-sectional estimation strategy feasible. For example, there are roughly ninety-five observations per trading day in Jonathan's sample, and although these are typically split across three or four different expirations, that still leaves about twenty-four to thirty-two different observations per expiration per day. In comparison, Bliss and Panigirtzoglou (2002) had 10.8 observations per day for their short sterling options data set, and found that their cross-sectional estimation strategy worked very well.

In the Bliss-Panigirtzoglou (2002) approach, the options prices on a given day are first converted into implied volatilities.[4] (Note that Jonathan does this same transformation in the run-up to equation (1); the idea here is exactly analogous to converting bond prices into bond yields to maturity—for most questions, it is more natural to work in "bond yield space" than in "bond price space," and in the options market it is likewise more natural to work in "implied volatility space" than in "options price space.") The prices of these different strikes, when converted into implied volatilities, typically form a "volatility smile," which is a smooth functional form that

[4]This is done by using the Black-Scholes formula assuming log-normality. As noted in the main text above, this is simply a way of transforming the options price data into "implied volatility space" rather than "options price space," because the former lends itself more naturally to fitting a smooth parametric function. Once this curve fitting is done, then the fitted yields can be transformed back to options price space. Thus, the log-normality assumption here plays no role in the analysis; it's simply a convenient way of transforming the data temporarily.

is easy to summarize with a small number of parameters. It only takes about five to ten such observations to get a reasonably good estimate of the smile (Bliss and Panigirtzoglou 2002), so the twenty-four to thirty-two observations in Jonathan's sample should be more than sufficient. Once the parametric functional form for the volatility smile is estimated, we can then transform it back to options price space and from there compute the implied PDFs using the second derivative of the options price with respect to the strike, as in Breeden and Litzenberger (1978).

Jonathan's non-parametric approach (following Ait-Sahalia and Lo 1998) relaxes the parametric functional form assumptions for the volatility smile required by Bliss and Panigirtzoglou (2002), but only by imposing alternative restrictions across time. This is because there aren't enough observations on any one day to get a good non-parametric estimate, so one must pool together observations on "similar days" and perform the non-parametric estimation on the pooled data. In this case, "similar days" are defined as those that have similar values for the assumed state vector Z, and Jonathan considers two different choices for Z.

Intuitively, for the non-parametric approach to be more accurate than the parametric approach, it must be the case that the intertemporal restrictions of the former are more consistent with the data than are the functional form restrictions on the volatility smile imposed by the latter. In my view, this is unlikely to be the case over long samples, such as the full 2009–16 pooled sample that Jonathan used in the first draft of this paper. Volatility smiles are, almost by definition, well approximated by a low-dimensional polynomial. Thus, I am very comfortable making parametric functional form assumptions here, and these restrictions are likely to fit the data quite well.

In contrast, I would be much less comfortable pooling together data from dates that may be years apart, as would be the case in the non-parametric approach applied to Jonathan's whole sample. In Jonathan's state vector S1, the "state of the world" is essentially completely described by the spot real interest rate r. (The other two state variables in S1 are the strike price and the maturity of the option, which one can think of as parameters of the option rather than the state of the world.) Thus, if the five-year TIPS yield was about 0.5 percent in November 2009, and is also about 0.5 percent currently, the pooled non-parametric approach would say that the data from November 2009 and today should be pooled together, because the state of the world is the same. This restriction strikes me as implausible, for many reasons; for example, the U.S. economy was close to the trough of a deep recession in November 2009, with the unemployment rate at 10 percent and the Federal Reserve pursuing unconventional monetary policy, while today the United States is in the midst of an ongoing expansion, with an unemployment rate of 5 percent and monetary policy that is acting much more conventionally. The true state of the world seems very different today than it was in November 2009.

Jonathan's second state vector, S2, provides a better alternative to S1 but still suffers from the same type of problem. In S2, the state of the world is completely described by σ_{EGARCH}, the estimated volatility of real yields rather than the level of those yields. (Again, there are two additional state variables in S2, the strike/spot price ratio and the maturity of the option, which describe parameters of the option rather than the state of the world.) The advantage of this approach is that it controls for uncertainty in the economy and only pools together days on which economic uncertainty was similar. Nevertheless, this restriction still seems much harder to justify than a smooth parametric functional form for the volatility smile.

In the revised version of the paper, Jonathan has greatly reduced these problems by performing the non-parametric analysis on each month in the sample separately. Thus, Jonathan's non-parametric approach no longer pools observations that are more than thirty days apart, so the state of the world across the pooled observations is almost certainly more similar than in the original, full-sample analysis. To the extent that the state of the world does change over the course of any given month, it may be well captured by the change in the real interest rate or the change in interest rate volatility. If so, then Jonathan's non-parametric approach should work just fine.

Nevertheless, one might worry that the state of the world changed rapidly in some months of the sample—particularly during the financial crisis and deep recession in 2007–9—and that the differences across days in some of those months are not well captured by a single state variable such as the real interest rate or interest rate volatility. For example, are we comfortable pooling observations at the beginning and end of September 2008, which are separated by the collapse of Lehman Brothers and AIG? And are we comfortable pooling observations at the beginning and end of March 2009, which are separated by the Federal Reserve's dramatic "QE1" announcement? Thus the same concerns I raise above could still apply, albeit to a lesser extent.

These concerns do not invalidate the non-parametric approach in Jonathan's paper, but they do raise legitimate caveats regarding that approach. A more parametric approach, along the lines of Bliss and Panigirtzoglou (2002) discussed above, could be used to produce snapshots of the options market on each and every day, which could provide a more accurate picture.

4. PDF Forecast Errors

Finally, Jonathan evaluates his estimated market-implied PDFs using the methods of Diebold, Gunther, and Tay (1998) and Berkowitz (2001). These methods apply the inverse of the cumulative distribution function (CDF) estimated at each point in time to map the outcomes of the predicted variable back onto the unit interval. If the model is doing a good job of producing PDFs, then the outcomes should be mapped back to a uniform distribution over the unit interval.

Figure 8 of Jonathan's paper shows that these mapped outcomes are quite far from uniformly distributed. But not only is this distribution not uniform, it is also heavily biased downward. This implies that the outcomes of the real interest rate between 2009 and 2015 were typically *lower* than the mean and median of the options-implied PDFs, and often fell in the first and second quintiles of that PDF distribution.

Thus, it is obvious from figure 8 that the options-implied PDFs were biased upward over this period. Jonathan takes this bias at face value and analyzes what it implies about the risk-neutral versus actual (or physical) probability measures underlying the options data. However, it's also possible that these under-realizations of the real interest rate are a feature of the sample period considered and would not be reproduced in a different sample. In support of this possibility, consider the evidence in Swanson and Williams (2014), who showed that bond yields, interest rate futures, and surveys all consistently predicted an increase in the federal funds rate throughout much of this period from 2009 to 2015, even though those federal funds rate increases never materialized. Thus, Treasury bonds, interest rate futures, and surveys from this period *all* tended to over-predict the nominal interest rate, and would also produce a transformed forecast error distribution that would look a lot like Jonathan's figure 8. Thus, it's not clear that those biased forecast errors in the TIPS options market should be taken as anything other than an idiosyncratic feature of the very special 2009–15 period.

5. Conclusions

In summary, Jonathan's paper introduces us to a new options data set and provides us with the first market-implied risk-neutral PDFs for the U.S. real interest rate (as opposed to the nominal interest rate or inflation). This kind of information is difficult or impossible to get any other way. Jonathan's non-parametric estimation strategy has been favored by some prominent authors (Ait-Sahalia and Lo 1998), but the parametric approach discussed by Bliss and Panigirtzoglou (2002) provides an attractive alternative that could perhaps produce more accurate estimates, with standard errors that are representative of a single day's data rather than a larger pooled sample. Although the options-implied PDFs that Jonathan estimates were biased over the 2009–15 sample, other interest rate forecasts were biased over that period as well, so one should hesitate before attributing that bias to TIPS option investors' risk aversion.

References

Ait-Sahalia, Y., and A. Lo. 1998. "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices." *Journal of Finance* 53 (2): 499–547.

Bank for International Settlements. 1999. *Estimating and Interpreting Probability Density Functions.* Proceedings of a workshop held at the Bank for International Settlements, Basel, Switzerland, June 14.

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.

Breeden, D., and R. Litzenberger. 1978. "Prices of State-Contingent Claims Implicit in Options Prices." *Journal of Business* 51 (4): 621–51.

Diebold, F., T. Gunther, and A. Tay. 1998. "Evaluating Density Forecasts with Applications to Financial Risk Management." *International Economic Review* 39 (4): 863–83.

Kitsul, Y., and J. Wright. 2013. "The Economics of Options-Implied Probability Density Functions." *Journal of Financial Economics* 110 (3): 696–711.

Swanson, E., 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.