Discussion

Richard Black

As is evident from the last two papers, the use of data from options prices is rapidly becoming a key tool for monetary policy.

Using the techniques described in both Levin, McManus, and Watt and Melick and Thomas, policy-makers can quickly estimate what the market is thinking either prior to or following a policy announcement or other change in information.

Estimates of risk-neutral density (RND) functions from options prices have many attractive properties:

- they are quick to calculate;
- they are readily understood and communicated (for example, the markets assessment of the probability of a 25 basis point increase in rates is x per cent);
- perhaps most importantly, they represent the market putting its money where its mouth is—a credibility-enhancing practice if there ever was one!

These attractive properties have certainly lead to the rapid use of RND functions in monetary policy. It is key then, to ensure that we have a full and thorough knowledge of the properties of the estimates, and it is here that the paper by Melick and Thomas makes, in my opinion, a welcome contribution.

In their paper, Melick and Thomas have taken two of the things that people frequently do (or at least what people frequently want to do) with RND functions and have analyzed two potential sources of problems.

The first thing that people often want to do is calculate a historical series corresponding to some attribute of the RND; for example, the skewness (or “bearish/bullish” coefficient) of the distribution or the kurtosis. In doing this, the problem of maturity dependence is often encountered. Melick and Thomas present a technique for addressing this problem.

Another thing that is done, or at least should be done, is attempt to assess the uncertainty surrounding the RND and any of the functions that are derived from it. Melick and Thomas have made a valuable start in addressing this problem.

In my comments, I will for the most part address the technical side of the paper, as the more market side was well commented on by the previous discussant. First, however, I would extend the comments by Melick and Thomas regarding the problematic issue of using an RND function to try to estimate market beliefs.

As noted by Melick and Thomas, and others, using RND functions does not necessarily, and in general does not, lead to a distribution that is useful for estimating actual market outcomes. I would like to provide an example of this. In the figure I have drawn two density functions for some variable of interest f at time T. The one on the right is the market distribution. Here the expected value if $E_P(f_T)$. The one on the left is RND and by the construction of an RND function the expected value is $e^{rT} f_t$.

Notice that there is no obvious relationship between the expected values of the two distributions, nor any other properties like variance, skewness, or kurtosis.

I have, of course, made the problem seem worse than it probably is in practice, where the two distributions are likely to be close. Nevertheless, I believe it an important caveat to be borne in mind.

Melick and Thomas go on the discuss how to derive constant-maturity RND functions. The issue here is that if underlying instruments have a maturity component to them, then the RND function will have one also. They note that there are two ways to handle the problem to either explicitly account for maturity in the RND function or calculate an ex post adjustment. In assessing the latter technique, they make innovative use of a constant-maturity FX option to gauge their success.

Melick and Thomas try a number of techniques to purge the RND functions, more specifically the interquartile range (IQR) implied by the RND functions, of maturity dependence and find that one specification in particular proved better than the others when compared with the OTC data. Although I found this general technique interesting, I thought that the section could benefit from the following:

- the authors should quantify the extent of the problem—for example, what is the difference between the IQR based on the nearby series from the CME data and the IQR based on the OTC data;
- assuming that the extent of the problem warrants taking into account maturity dependence, the authors should describe exactly how to use the "detrended" IQR to estimate the actual interquartile range;
- for the practitioners who do not have an OTC-equivalent data set, the authors should describe how well traditional measures of fit (for example R^2) can be used to gauge which detrending method works best;
- finally, it would be useful if this technique was compared with that where maturity dependence is built into the RND function directly. This would be a valuable piece of information for anyone about to make some adjustment for maturity dependence.

In the next section of the paper, Melick and Thomas go on to discuss and illustrate methods for gauging the uncertainty surrounding the estimates of the RND functions and estimates derived from them.

In their paper, the authors outline how they view the problem. Essentially it is a problem of curve-fitting with an assumption that deviations from the fitted curve are random. By taking great liberties one could illustrate this as follows.

In Figure 1 the error bars describe the pricing errors, which the RND function is chosen to minimize. Melick and Thomas then go on to try and assess the uncertainty as represented by these errors.

Before going on to consider these methods, it is worth noting that there are some potential problems with this approach. In particular, by choosing a functional form flexible enough one could eliminate all errors without any formal penalty. This means it is not obvious that it is sensible to think of these errors as coming from a distribution as they may be generated simply as a by-product of the functional form of the fitted density.

That said, this is a fundamental problem in this area, and the approach taken by the authors is, in my opinion, very reasonable.

The authors propose two methods for estimating the error variance. In the first method, the authors utilize their assumption that the errors are random and that the coefficients are estimated by maximum likelihood to calculate the asymtotic distribution of the coefficients. Using this distribution, different sets of coefficients for the RND are drawn and using these an error distribution is built up. This method, as the authors noted, has several drawbacks. Among them is the assumption that the asymtotic distribution has been reached—an assumption that, given there are only 23 data points, is rather strong. Another assumption is that the error terms are

Figure 1

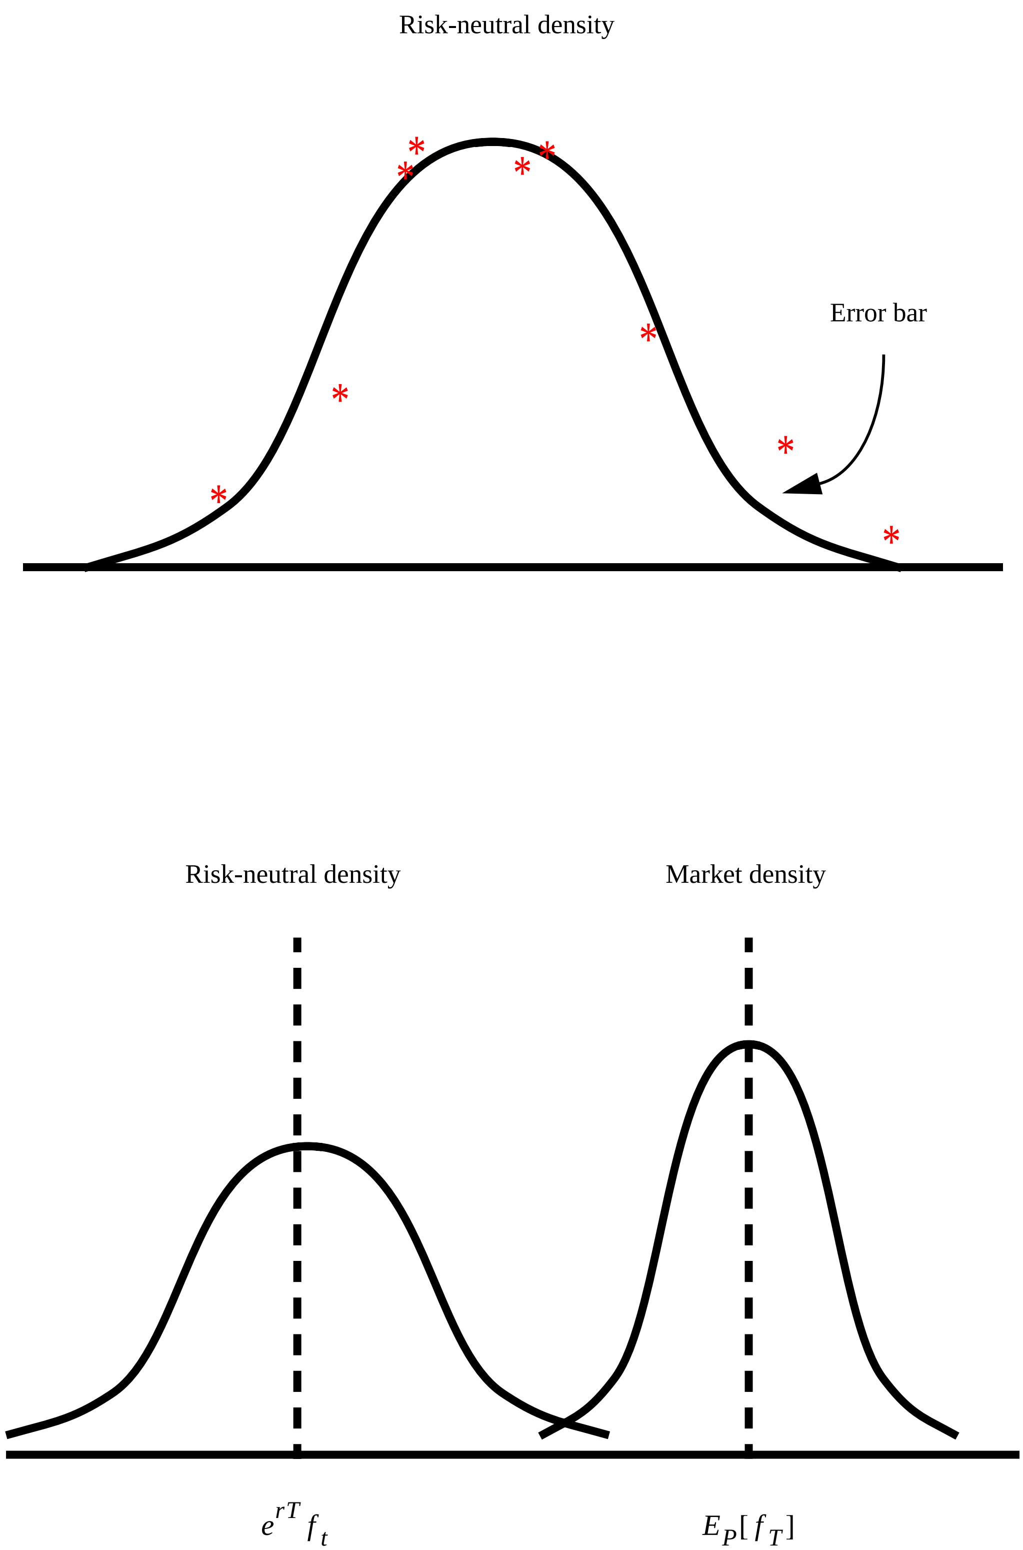
Risk-neutral density
Error bar
Risk-neutral density
Market density
$e^{rT} f_t$
$E_P[f_T]$

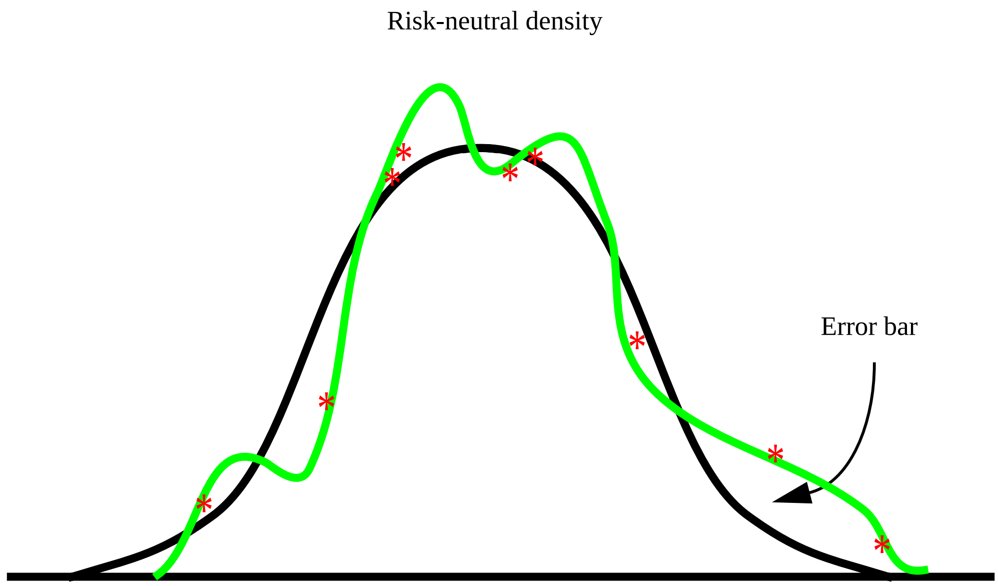

independent. Melick and Thomas discuss this point and note that the pricing errors are not independent.

In their second approach, pseudo-samples are created from the original data and the RND function is recalculated. In this way the distributional assumptions made in the Monte Carlo method are avoided. The problem here, as noted by Melick and Thomas, is that the special nature of the data set makes such pseudo-samples fundamentally different from the observations.

One thing that crossed my mind is whether the authors had tried a method that may be described as "error" Monte Carlo. This method aims to combine the advantages of the Monte Carlo and bootstrap methods.

This method is similar to the bootstrap method, except that instead of drawing the pseudo-sample from the observed data set, it draws new shock terms from the implied pricing errors and uses these to generate a new sample. In this way, the data order is preserved and less distributional assumptions are required compared with the Monte Carlo method. Nevertheless, the residuals should still be independent, something which they plainly are not. The only thing I could suggest is that the authors try to model the time series properties of the residuals and to use the results to purge the residuals of any correlation.

In summary, the paper by Melick and Thomas provides a valuable starting point from which to develop measures of accuracy for RND functions. As these methods are developed they will, I hope, become more prevalent and replace the eyeball metric that is so often used currently.