Methodology for Estimating Risk Neutral Probability Density Functions

We estimate risk neutral probability density functions (RNPDs) for a variety of different asset classes using a variation of the technique developed by Shimko (1993). This procedure involves fitting a curve to the implied volatilities of a series of options and expressing the volatility as a function of the strike price. The implied volatilities are then translated into continuous call option prices, and the risk neutral distribution of the underlying asset is obtained through the Breeden-Litzenberger (1978) method.

Shimko Procedure

The data we use for the Shimko procedure are based on actual trades that occurred, as opposed to simple bid/ask quotes from dealers. In order to implement the procedure, we need a sufficient number of option price observations to successfully fit the volatility curve. Moreover, the observations also need to span a sizable portion of the available strike price range so that we can minimize the amount of the corresponding density function that is based on extrapolation instead of empirical observations. As such, we utilize prices from option trades that occurred over a five-day window. While many option markets, like those for agricultural commodities such as wheat and corn, typically have 30 to 40 transactions that occur on a single day at different strike prices, other option markets, such as those on the one-month crude oil futures contract, trade more sporadically. Indeed, in July 2010 there were, on average, only eight different strike prices with any activity in the one-month oil futures option market.

We are interested in generating the probability density function for a specific asset as of the current date and then comparing it to the density functions that existed one month earlier. To that end, we perform the following steps (this description will focus on options with the S&P 500 Index as the underlying asset, and it assumes the current date is 9/30/2010):

1. We identify the option that expires approximately six months from the current date. For options that are issued quarterly, this means that we will be dealing with expirations that range from five months to seven months. Once the correct option has been identified, we obtain from Bloomberg the closing prices for that option across all strike prices that traded during the five-day window that ends on 9/30/2010. Note that we use both call and put options. Since Bloomberg often provides end-of-day settlement prices in certain markets when no trading has occurred, we require that the corresponding volume in each option on the day be greater than zero. There were trades in 35 different strike prices in this five-day period, ranging from 400 to 1450, for a total of 120 observations (the same strike price could trade multiple times during the period).

2. We calculate the time to expiration of each option as the number of days between the price date and the option expiration date based on a 30/360 calendar (12 30-day months), divided by 360.

3. We associate a risk-free rate with each option based on an interpolation of the U.S. Treasury yield curve (consisting of on-the-run issues) that existed on the specific price date. We linearly interpolate the yield curve to match the time to expiration of each option and express each value as a continuously compounded interest rate.

4. All put options are transformed into call options using the put-call parity appropriate for the asset class. In the case of options on the S&P 500 Index, we use the following:

$$C = P + Se^{-\delta\tau} - Xe^{-r\tau}$$

where C is the call price, P is the put price, S is the spot price of the underlying asset, X is the strike price, r is the continuously compounded interest rate, δ is the continuously compounded dividend yield, and τ is the time to expiration.

5. We calculate the implied volatility, σ, of each option using an option pricing model appropriate for the asset class. In the case of options on the S&P 500 Index, we use the standard Black-Scholes model:

$$C = Se^{-\delta\tau}\Phi[d_1] - \mathrm{X}e^{-r\tau}\Phi[d_2]$$

$$d_1 = \frac{ln\left(\frac{S}{X}\right) + \left(r - \delta + \frac{\sigma^2}{2}\right)\tau}{\sigma\sqrt{\tau}}$$

$$d_2 = d_1 - \sigma\sqrt{\tau}$$

where $\Phi(.)$ is the cumulative normal distribution, and solve for the implied volatility. We eliminate all observations where the call option price does not permit a volatility to be calculated.

6. Next, we condense the data set by selecting a single option at each strike price. If multiple observations are present, we choose the most recent price date and give primacy to call options over transformed put options if the strike price is greater than the current spot price (when the strike price is below the current spot price we give primacy to transformed put options). In essence, we treat the five-day period as if it were a single trading day and select the "last price" at each strike.

7. We then fit a cubic B-spline smooth to the [strike price, implied volatility] pairs, interpolate several hundred implied volatilities between the minimum and maximum strike prices, and use the interpolated values to calculate call option prices using the model specified in step 5. Note that we use as inputs to the model the value of the spot price, risk-free rate, dividend yield, and time to expiration at the end of the five-day window. We use the first and last estimated spline segments to extrapolate beyond the minimum and maximum strike prices present in our data set. The resulting [strike price, call price] observations serve as our proxy for the continuous call price function.

8. The first and second derivatives of the proxy call price function are estimated numerically, and the values of the second derivative are divided by $e^{-r\tau}$. The result is the risk neutral probability density function (RNPD) implied by the options market for the future value of the S&P 500 Index as of 9/30/2010.

9. Lastly, we convert the RNPD from "price space" into "log return space" via the standard change-of-variable transformation.

For simplicity, we categorize all of the RNPDs—regardless of the actual time to expiration of the option we used—as being "six-month" density functions and interpret our results to be the distributions of the expected values six months from each observation date. We utilize the Black option pricing formulas for contracts based on futures contracts (which apply to commodities like gold, oil, and corn) and the Garman-Kohlhagan pricing formulas for options based on exchange rates.

References

Breeden, D. T., and Litzenberger, R. H. (1978), "Prices of state-contingent claims implicit in option prices," *Journal of Business* 51 (4), pp. 621-51.

Shimko, D. C. (1993), "Bounds of Probability," *Risk*, 6 (4), pp. 33-37.