Revealing Market Expectations with Derivatives

Bruce Grundy, Finance, UoM, Professor of Finance;
Paul Kofman, Finance, UoM, Professor of Finance

1. Background and aims of project

A fundamental function of financial markets is price discovery. Instantaneous access to, and continuous dissemination of, traded prices are necessary conditions for a financial market to effectively perform its price discovery function. Whereas many cash markets fail these conditions, their derivatives markets often act as substitutes for price discovery. Good examples of such cash market ‘failure’ can be found in commodity markets, where the cash markets tend to be opaque (e.g., for crude oil or coffee). While cash market transparency has undoubtedly been improved by internet technology allowing real-time disclosure of global market information, cash traders in those commodities will still routinely benchmark against derivatives prices.

Of course, derivatives prices are distinct from cash prices as they reflect demand/supply conditions at different points in time. So, what information is exactly revealed by derivatives prices? Do derivatives prices reveal how the market thinks cash prices will evolve over time? Or do they reflect something altogether distinct, ‘disconnected’ from cash prices? Finance theory establishes that convergence and no-arbitrage invoke strong relationships between cash and derivatives prices. The cost-of-carry and the Black-Scholes pricing models are key examples of such relationships. Underlying these relationships is the notion of fair value based on the conditional expectation of future value of the payoffs. This conditional expectation is a function of the underlying probability distribution, and its computation requires the input of the parameters characterizing this distribution. Of course, rather than imputing parameter values to solve for fair value, one can also take market prices as fair value and instead solve for the market-implied parameters by inverting the derivative asset pricing equation..

Option prices, futures/forward prices and other derivatives prices allow us to infer ‘market expectations’ on a variety of (related) distribution metrics. While option implied volatility is undoubtedly the best known metric, other measures of interest might include higher order moments like implied skewness or implied kurtosis. Less ‘statistical’ measures include implied risk premia, and implied risk aversion. Increasing availability of traded derivatives prices has made it feasible to extract a large array of market expectations, yet the interpretation and implications of these expectations has yet to be fully appreciated by academics and practitioners alike. This project can therefore be expected to contribute significantly to the finance literature, as well as it can be expected to deliver recommendations for the development of new derivatives markets of special interest to Australia.

The aims of this ‘umbrella’ project are summarized as follows. The applicants will investigate a range of (semi-)parametric and nonparametric methods to recover both the implied risk-neutral distribution as well as the ‘true/objective’ distribution of the underlying asset from market prices of various derivatives. In particular, the suitability and ‘accuracy’ of these methods will be assessed empirically for a range of traditional and novel financial markets, including those trading:

- commodity derivatives
- currency derivatives
- energy derivatives
- economic derivatives
- equity derivatives

- fixed income (CDO/CDS) derivatives
- insurance (CATs) derivatives
- real estate derivatives (incl. MBS)

For each of these markets – where appropriate – we will characterize the underlying distribution by inferring the implied

a) volatility expectations/processes/term structures;
b) correlation (term)structures;
c) risk attitude measures;
d) hedge parameters.

2. Significance and innovation

Despite the fact that the concept of implied distributional parameters was first discussed soon after the publication of the Black-Scholes formula (see e.g., Latane and Rendleman, 1976 and Breeden and Litzenberger, 1978), the following two decades saw remarkably little development in deriving the broader implications of traded option prices for inference on market expectations regarding the underlying assets. Only with increased access to traded option prices did academics pursue more challenging quests for derivatives-implied market insights. At the same time, global financial market deregulation required (and allowed) new markets in ever more sophisticated underlying assets where traditionally 'price' discovery had been obscured by the absence of (continuous) information flows and the nature of the often dispersed cash markets. This revived academic interest in the 90s – reviewed by Jackwerth (1999) – focused predominantly on the technology to recover the distributional parameters of interest from traded option prices. The following section will briefly summarize the status quo and address certain critical issues.

While the 'technical' papers certainly illustrated their methodology for particular markets, the more specific examples tend to be restricted to events like the implied oil price expectations during the Gulf crisis by Melick and Thomas (1997), or the implied futures prices during limit moves by Hall, Kofman and Manaster (2006). This project will consider the empirical implications in much greater detail and provide significant coverage of markets, existing and newly developing. In particular, we will consider markets still lacking but of potential great interest and significance to Australia. Applications to each of these markets will require tailored technologies, but we will start with an 'umbrella' methodology, and then investigate specific methods for density inference calibrated to the relevant market conditions. This will allow us to assess whether the hypotheses to be tested are dependent on the choice of density estimation, rather than directly test and compare the different methodologies. We will also investigate new methods to infer the 'true' probability density from derivatives-implied risk neutral densities.

There is considerable practical as well as theoretical value and interest in these implied measures. Investors frequently resort to implied parameters to help them value related investment opportunities. Rating agencies use implied default probabilities to signal rating changes. Policy makers could (and do) use combinations of these measures to assess whether their policy intentions are properly assessed by the market. The Bank of England, for example, is one of the more prominent users of the derivatives implied expectations technology to benchmark its policy directions/actions against market expectations. Theoretical interest is mainly focused on the implications of risk-neutrality and implied risk attitude. We intend to investigate different ways to present and interpret the implied distributional metrics tailored to the specific needs of end-users, both practitioners and academics.

3. Description of Approach

The following projects share a common theme: the use of traded derivatives prices to extract market expectations. The combined projects will investigate both parametric and nonparametric methods to recover risk-neutral densities from market prices of the various derivative assets. While the basic methodology will be identical, each specific project will require tailored functional specifications to match market characteristics.

The technology to extract market expectations has made significant progress in the past 15 years. Basically, the literature distinguishes two methods. The first method uses the traded option prices to provide identifying conditions by either choosing model parameters or, by calibrating a binomial tree or volatility function to match observed option prices. This method has been extensively described in Jackwerth and Rubinstein (1996) and Jackwerth (1999). The second method estimates the second derivative of the option price function relative to the strike price as the implied (risk-neutral) density. Breeden and Litzenberger (1978) already demonstrated this, but the technology is further developed in Shimko (1991).

There are a few issues of concern with both methods. First, the model (say Black-Scholes) assumptions may not be appropriate. So-called model risk can be (and has been) considered and we will investigate the sensitivity of our results accordingly. Next, the traded option prices may be of poor ‘quality’ and not reflect fair value. An appropriate methodology matches the quality and quantity of option price data to the choice of technology. An abundance of high quality data allows non-parametric estimation, while limited quantity – lower quality data typically requires parametric estimation. The quantity/quality is then usually a function of the liquidity of the option price series. Access to high quality (and quantity) option price data will significantly enhance the reliability of our results. Finally, and most importantly, the implied density is by definition a risk-neutral density, not the ‘true’ density of the underlying asset. Whether the risk-neutral density provides useful information concerning the probabilities of future outcomes is still an important open question. Whereas the location parameter is almost certainly a biased estimate of the true expectation, empirical evidence seems to suggest that the scale parameter and higher order metrics are less affected when comparing realized densities with implied risk-neutral densities. A few academics have considered this discrepancy between risk-neutral implied density and the representative risk-averse subjective (true) density. Grundy (1991) derives conditions under which traded option prices can constrain the feasible set of distributional moments for the underlying asset to learn about the true distribution. Bliss and Panigirtzoglou (2004), on the other hand, combine various utility specifications with the implied risk neutral density to provide the best possible fit between the realized and subjective density. This approach provides the researcher with density forecasts, as well as a measure of implied risk attitude. The approach is based on the notion that a representative agent’s subjective density function can be obtained by dividing the risk-neutral density by the normalized marginal utility of the agent. The latter can then be parameterized according to common utility specifications, e.g. power utility. Anagnou, Bedendo, Hodges and Tompkins (2002) find that a utility-adjusted implied density does indeed seem to provide unbiased forecasts of the realized density unlike a variety of ‘raw’ risk-neutral densities (based on different estimation methods). Aït-Sahalia Wang and Yared (2001) avoid making assumptions (or estimating) representative preferences. Instead, they use Girsanov’s characterization of the change of measure from the actual density to the option-implied density where the diffusion function of the underlying dynamics is identical, but the drift needs to be adjusted. The drift adjustment can be inferred from other derivatives (futures) on the same underlying asset. Of course, that approach restricts the applicability and limits the interpretation of the implied density metrics.

Our extraction methodology will commence with a comprehensive intertemporal analysis of implied densities using traditional options markets. This will allow us to benchmark the robustness and efficiency of the different implied density estimation methodologies. Our analysis of forecast errors focuses on various density metrics, including location, scale/dispersion, asymmetry, fat-tailedness, relevant tail percentiles. We will then continue our analysis for 'new' markets, examples of which are given below for specific implied metrics of interest.

Implied Expectations

Following Breeden and Litzenberger (1978), Manaster and Rendleman (1982) used option prices to infer market predictions of equilibrium stock prices. In so-called prediction markets, payoffs are tied to the outcome of future events. The prices of prediction assets reveal the market's assessment of the likelihood of the event occurring. For example, the recent availability of traded *economic* derivatives prices at the Chicago Mercantile Exchange (the CME Auction markets) will significantly enhance our understanding of the market's assessment of possible monetary policy actions/outcomes. While the Federal Funds Rate Futures (see Easton and Pinder, 2007, for an Australian equivalent – the 30-day Interbank Cash Rate Futures) has been around for some time, it can only provide limited information on the market-implied probability of a single (direction) change in interest rates. The economic derivatives traded option prices, however, allow us to assess forecasts of a complete probability distribution of changes in say, interest rates (but also on GDP, trade balances, CPI). An issue worth investigating is whether market expectations are likely to be asymmetric around monetary policy actions. During monetary policy tightening, option-implied distributions are likely to be positively skewed, indicating that market participants attach higher probabilities to sharp interest rate increases while this asymmetry is reversed during loosening monetary policy. Market expectations could also vary significantly intertemporally, as asymmetries in market expectations tend to increase before changes in the monetary policy, and to decrease thereafter.

While a few academic papers have started examining the information content of these economic derivatives prices, none so far have considered the joint probability densities on (macro)economic outcomes. Yet another issue for investigation is the comparison of market-implied interest rate expectations based on the T-Bill interest rate derivatives prices (that have been around for 25 years) and the more 'focused' economic derivatives prices (based on target rates set by the monetary authorities). Bhar and Chiarella (2000), for example, have used the former (90-day BAB futures/options) to imply market expectations regarding the RBA target rate. It will be worthwhile to assess the robustness of their results in a US comparison of TBill futures/options implied expectations against the relevant economic derivatives implied expectations. The outcome of that comparison could suggest whether it is worthwhile to introduce similar derivatives in Australia.

Implied Risk (and Risk Attitude)

Trade in fixed income derivatives has experienced tremendous growth over the past decade. While trade in (and new issues of) treasury securities has significantly diminished, this market decline has been more than offset by booming trade in corporate securities. These securities carry significant additional risks, which explains the matching growth in traded CDOs and CDSs. For our project's purpose, traded CDO/CDS prices allow inference on market-implied default risk (premia). The relationship between default risk premia and default probability is not necessarily monotone as default risk premia could increase with increasing probability of default when this default probability is low, but decline when the default probability reaches a certain threshold level. A second topic worthy of investigation here, is the shape of the default risk premium term structure and its response to business cycle variations.

Another project within this risk category consists of market risk expectations as revealed by real estate derivatives prices (futures and options now trade on the Chicago Mercantile Exchange, while swaps trade in the UK). While this market in retail and commercial real estate derivatives has only just commenced in the US, its implications will be of particular interest for Australia. Some of the difficulties encountered in the proper design of the underlying indices would be of lesser concern in Australia with its geographically segmented real estate market structure. Inference on real estate market volatility as revealed by traded real estate derivatives prices could be of significant benefit to the RBA as a barometer of inflationary pressure.

Yet another market-implied risk application involves broadly defined commodity markets, including the energy and weather derivatives markets. Compared with equity/fixed income risk, commodity price risk is excessive and typically unstable. Energy prices and weather indices display even more significant distributional complexity, with extreme outliers and excessive skewness/kurtosis. Rather than the location parameter, the dispersion is the main measure of concern to commodity market participants. Extreme non-normality (and its impact on forecasting prices) stresses the need to better understand the intertemporal movements of risk-based density metrics. Intertemporal option-implied densities may well facilitate this learning process.

Implied Correlation

Following from implied risk, it is a small step to also infer market-implied correlations between sets of underlying assets. Multi-asset options, like foreign exchange options (aka quantos or more generally, rainbow options, basket options, exchange options) allow a characterization of the market-implied covariance matrix. Recent conditional correlation estimation methodologies include extreme value, truncated, exceedance and copula estimators. Our project will evaluate the relative merits of these technologies when compared with option-implied conditional correlations. Further analysis of implied correlation structures will also allow us to measure the existence of a correlation risk premium. Using the implied covariance matrix to forecast correlated risks (and correlation risk premia) is of particular relevance in a Value-at-Risk (VaR) framework. Implied correlation is also increasingly used by fixed income practitioners when comparing alternative investments in synthetic CDO tranches (but also for Mortgage Backed Securities). The estimation methodology for implied correlations needs to allow for the heterogeneity of the underlying portfolio to be of use in relative value analysis of alternative tranched investments. The advent of correlation trading underlines the importance of further developments in this line of research.

References

Aït-Sahalia, Y., Y. Wang and F. Yared, 2001. Do Option Markets Correctly Price the Probabilities of movement of the Underlying Asset? *Journal of Econometrics* 102, 67-110.

Anagnou, I., M. Bedendo, S.D. Hodges and R. Tompkins, 2002. The relation between implied and realised probability density functions. Working Paper, Financial Options Research Centre, University of Warwick.

Bhar, R., and C. Chiarella, 2000. Expectations of monetary policy in Australia implied by the probability distribution of interest rate derivatives. *European Journal of Finance* 6, 113-125.

Bliss, R. and N. Panigirtzoglou, 2004. Option-Implied Risk Aversion Estimates. *Journal of Finance* 54, 407-446.

Breeden, D. and R. Litzenberger, 1978. Prices of State-Contingent Claims Implicit in Options Prices. *Journal of Business* 51, 621-651.

Easton, S. and S. Pinder, 2007. Predicting Reserve Bank of Australia Interest Rate Announcements Using the 30-day Interbank Cash Rate Futures Contract: Beware the Target Rate Tracker. *Australian Economic Review*.

Grundy, B.D., 1991. Option Prices and the Underlying Asset's Return Distribution. *Journal of Finance* 46, 1045-1069.

Hall, A.D., P. Kofman and S. Manaster, 2006. Migration of Price Discovery in Semiregulated Derivatives Markets. *Journal of Futures Markets* 26, 209-242.

Jackwerth, J.C., 1999. Option-Implied Risk-Neutral Distributions and Implied Binomial Trees: A Literature Review. *Journal of Derivatives* Winter, 66-82.

Jackwerth, J.C. and M. Rubinstein 1996. Recovering Probability Distributions from Option Prices. *Journal of Finance* 51, 1611-1631.

Latané, H. and R.J. Rendleman, 1976. Standard Deviation of Stock Price Ratios Implied in Option Prices. *Journal of Finance* 31, 369-381.

Manaster, S. and R.J. Rendleman, 1982. Option Prices as Predictors of Equilibrium Stock Prices. *Journal of Finance* 37, 1043-1057.

Melick, W.R. and C.P. Thomas, 1997. Recovering an asset's implied pdf from option prices: An application to crude oil during the Gulf crisis. *Journal of Financial and Quantitative Analysis* 32, 91-115.

Shimko, D., 1991. Beyond Implied Volatility: Probability Distributions and Hedge Ratios Implied by Option Prices. Working Paper, University of Southern California.