THE JOURNAL OF DERIVATIVES

VOLUME 13, NUMBER 2 WINTER 2005

One of the most important issues in applying contingent claims models derived from theory to trading actual derivative instruments in the real world is the imperfect connection between empirical or "true" probability distributions and "risk neutral" probability distributions. True probabilities govern the occurrence of events in the real world. Risk neutral probabilities apply to the pricing of payoffs that are contingent on those events. True probabilities can be estimated statistically from event frequencies in the past and other historical data. Risk neutral probabilities are implied out from market prices for options and other derivatives using a pricing model.

One of the most familiar examples of the dichotomy between the two is the volatility smile. A given stock can have only one (true) volatility at a time, but the (risk neutral) implied volatilities (IVs) from a set of options all based on that stock will typically exhibit a range of values, with out of the money puts having distinctly higher IVs than at the money options. Is this evidence of gross mispricing in the options market and persistent arbitrage opportunities? Or is it evidence that the model used to obtain the IVs is incorrect, perhaps because investors know that the true distribution has fatter tails than the one the model is derived from, so they price options systematically differently from what the model dictates. In that case, the it is the market that is right and the model is wrong. Or does the smile, which turned into a persistent "smirk" after 1987, just reflect a kind of aversion to risk—"Crash-o-phobia," as it has been called—such that investors don't disagree with the model, but are nevertheless willing to pay more for put protection than what it is worth in actuarial terms, simply to obtain peace of mind? The answer is. . . that we don't know the answer.

Our lead article doesn't resolve the dichotomy between true and risk neutral probabilities either, unfortunately, but it does provide a large amount of new information about the volatility smile in equity markets around the world. Foresi and Wu study a comprehensive high quality set of IV data from 12 countries and find strong similarities among them. Every market displays a volatility smirk, and they all seem to get stronger for longer maturity contracts. Changes in the levels and slopes show strong correlation across the different markets. So while we end up still unsure of its origins, it is clear that the volatility smirk in implied volatilities from equity index options is a worldwide phenomenon.

THE JOURNAL OF
DERIVATIVES

Our second article deals with an important aspect of practical interest rate modeling: calibration. The LIBOR Market Model is widely used in practice, but it has a lot of parameters to fit and they have to obey certain constraints. The Markov-functional model is an efficient way to do this. Bennett and Kennedy show that it is also a very general way, because any accurate approximation to a separable LIBOR Market Model will also approximate a Markov-functional model. Another case of straightforward theory that can lead to difficulties in implementation is pricing options on a dividend paying stock. Subtracting the present value of future dividends from the current stock price in the pricing equation is fine for short maturity options when there are only one or two highly predictable payouts before expiration. But for long-dated options, payouts must really be treated as stochastic, and not perfectly predictable. Korn and Rogers resolve this difficulty by modeling the stock price as a function of the underlying dividend process, bringing dividend uncertainty directly into the framework of the model.

Our last two articles are concerned with issues of hedge design. Nam, Tucker and Wei confront the risk management problem faced by farmers and other commodity producers who face both price and quantity risk, the latter being comprised of idiosyncratic and aggregate exposures. They demonstrate the efficacy of an option whose payoff is based on the commodity price, but with a trigger based on quantity, as a hedging vehicle for this case of combined risk. Finally, Grieves and Marcus point out a problem in hedge design created by the "quality option" in a bond futures contract: When the Cheapest to Deliver (CTD) bond changes, it will often flip to a new CTD bond with quite different hedging characteristics. The right approach is to use a hedge ratio that takes the delivery option into account explicitly.

By the time you are reading this, we will entering the year-end holiday season. Not to mention the bonus season. We wish you much happiness on all fronts.

Stephen Figlewski
Editor