DAVID BATES ON CRASH AND JUMPS

David Bates is a Professor of Finance at the University of Iowa. He early on specialized in jumps and stochastic volatility. David Bates was one of the first to show how implied distributions from liquid options could be used to extract useful market expectations. He has shown us that options not only are of interest for traders and hedgers but also for anyone that wants to extract information about what some of the smartest people in the market (option traders) know and expect.

David Bates is not only developing fancy models he is an empiricist, spending considerable time testing out if the models really work reasonable well. As an option trader I am particular interested in tail events, and particular what I call models of models risk. When most people buy an exchange traded put option and the underlying stock falls in value below the strike price they think for certain they will get paid. Most people forget that underlying their derivatives model they have a model of default risk in the clearinghouse. David Bates has empirically looked into such types of risk something we soon will learn more about:

Haug: When did you first get interested in quantitative finance?

> Largely by accident. As a grad student at Princeton, my primary interests were in international finance. I was trying to come up with an explanation for uncovered interest parity rejections – the fact that putting money in relatively high-interest currencies has been surprisingly profitable. I thought a peso problem might be the answer, and currency options struck me as the right place to look for evidence of peso problems. Once I started, I got more and more interested in options.

Haug: What exactly is a peso problem?

> A "peso problem" is a rare event: something major that investors think may occur, but that econometricians have trouble identifying statistically. The term refers to the Mexican peso in the 1970's. Mexican-U.S. interest differentials suggested a future peso devaluation, yet the exchange rate remained fixed for years – until the peso was ultimately devalued in 1976. I was hypothesizing that something similar might be happening with the dollar versus European currencies in the early 1980's: investors might have been deterred from investing in relatively high-interest dollars by fears of a dollar crash. I figured such crash fears should show up in high prices of out-of-the-money call options on Deutsche marks – which I did indeed find in 1984 and early 1985, as the dollar peaked.

Haug: What is your educational background?

> I was an undergraduate math major at MIT. I loved math, but I realized there that it would be hard for me to push the frontiers. So I gradually switched to economics, ultimately ending with a Ph.D in economics from Princeton. Princeton was a great place at the time to do finance; John Campbell, Pete Kyle, and Sandy Grossman were there.

Haug: You were to my knowledge one of the very first to extract implied distributions from liquid option prices in a paper you published in Journal of Finance in 1991. How did you come up with this idea and what was the paper about?

> It was really a spin-off of the peso problem research in my dissertation. I was looking at foreign currency options for evidence of fears of a dollar crash – this was in the summer of 1987. When the stock market crashed that fall, it was obvious that the same methodology could be used with stock index options. So I got CME data for options on S&P 500 futures, and ran the same diagnostics I had run for Deutsche mark futures options: my skewness premium measure of implicit skewness, and daily implicit parameter estimates of a jump-diffusion model roughly based on Merton (1976). Backing out implicit jump-diffusion parameters was essentially the same idea as computing implicit volatility; but the more general model allowed greater flexibility in implicit distributions and assessing crash risk.

Haug: Based on the implied distributions you backed out did the stock option market anticipate the stock crash of 1987?

> Yes and no. I did find substantial fears of a crash implicit in option prices during the year before the actual crash; in particular, during October 1986 to February 1987 and again in June to August, 1987. Indeed, the negative implicit skewness during those periods was comparable to the pronounced negative skewness we have observed *since* the crash of ‘87. However, the implicit crash fears subsided quickly after the market peaked in August 1987, and were negligible in the weeks leading up to the crash.

Haug: What are the difference and similarities between implied risk neutral distributions and expected real distributions, is it only the mean that is different, what about volatility, skewness, kurtosis and even higher order moments?

> The conditional mean always changes. The impact on other moments depends upon what sorts of risks are being priced: stochastic volatility risk, for instance, or jump risk. A stochastic volatility risk premium primarily affects the term structure of expected average variances (actual versus risk-neutral), but has less impact on other moments. At short maturities, the actual and risk-neutral variances should be identical if the only risk is stochastic volatility.
>
> A jump risk premium primarily changes the actual versus risk-neutral jump intensity; there are also effects on the mean jump size. These directly affect variance assessments at *all* maturities, but also affect skewness and kurtosis. For stock market crash risk, the risk-neutral jump intensity is higher than the actual jump intensity, creating higher variance under the risk-neutral distribution and also more downside risk.

Haug: You where also the first one to describe put-call symmetry, how was this related to extracting implied distributions?

I called it the "skewness premium:" the proposition that call options x% out-of-the-money should be worth x% more than comparably OTM put options, for a broad class of quasi-symmetric distributions with small degrees of positive skewness. I had noticed it was a property of the Black-Scholes-Merton call and put formulas. I subsequently realized it was also a property of other popular option pricing models: the Hull and White (1987)/Scott (1987) stochastic volatility option pricing models, the Merton (1976) jump-diffusion model with mean-zero jumps. Furthermore, it held for exchange-traded American futures options, for which I could get data, as well as for European options in general.

The advantage of the skewness premium was that I could say *all* of these models were wrong, and that we needed models with more skewness, without having to get into model-specific estimation. It was a simple and intuitive nonparametric diagnostic. Now, of course, we routinely build skewness into our option pricing models.

Haug: How was the jump-diffusion model you introduced in 1991 different from the Merton-76 jump-diffusion model?

Primarily in pricing jump risk. Merton's model assumed jump risks were firm-specific and idiosyncratic, so that the actual and risk-neutral jump intensities and jump distributions were identical. Since I was working with stock index options, for which jump risk must be systematic, I had to come up with a method of pricing that risk. I did that in my dissertation, by extending the Cox, Ingersoll, and Ross (1985) general equilibrium production economy model to jump-diffusions. My 1991 paper did it more simply, by working more directly with the distribution of what we now call the "pricing kernel."

Haug: What are the fundamental reasons for jumps in financial markets?

Good question, but I don't know the answer. Small jumps can often be attributed to announcement effects, or major pieces of news. However, it's the large jumps that we worry about, and we don't really know what causes those.

Haug: If you use an option model assuming continuous hedging and you cannot hedge continuously in practice due to transaction costs etc. then I assume this also must have impacts on the implied distribution you are backing out from an option model assuming continuous hedging?

Not necessarily. It is possible to derive option pricing models from equilibrium grounds, rather than from no-arbitrage grounds. Indeed, the original Black-Scholes paper primarily derived the formula from the continuous-time CAPM. I think I recall that Fischer Black preferred that justification, precisely because he was skeptical as to whether the no-arbitrage strategy was implementable.

Haug: Every good option trader I am aware of know that jumps are extremely important for option valuation and especially out-of-the money options with short time to maturity. Still relatively few option traders I know use option models that takes into account jumps, instead they try to fudge a simpler model. Why is this, is it to difficult to estimate the parameters, are the jump models not good enough or is it simply most traders that not are smart enough?

> The analogy I draw is Newton versus Ptolemy in planetary orbits. Academics are interested in identifying the underlying deeper structure: the fact that orbits are elliptical. Practitioners are interested in what works; their fudge factors are the equivalent of epicycles and epicenters. And such fudge factors did work; I gather the Ptolemaic model could indeed predict planetary orbits as well as the Newtonian model, for the data available in Newton's time.
>
> In option pricing, we don't yet have a simple, perfectly fitting model equivalent to Newton's model. A simple jump model has difficulty matching option prices at all maturities, even when calibrated to one day's data. We're moving to more complicated models that do fit progressively better; but maybe we too are just engaging in the equivalent of adding more epicycles and epicenters. Furthermore, the academic research is focused more on pricing, whereas practitioners are more interested in hedging. So it is perhaps not surprising that practitioners have their own methods, which presumably work for them.

Haug: Even if I calibrate a jump-diffusion model to many years of historical data I would think I still always will underestimate the potential jump size, even if I include the crash of 87 in my data how do we know that we will not have an even bigger crash at some point ahead?

> I currently view the '87 crash (down 23%) as a "normal" 10-15% crash that was made worse by the DOT system getting overloaded; remember that the market rebounded the following two days. If the order flow system has adequate capacity, then another '87-like crash would appear unlikely. But predicting rare events is inherently difficult, as illustrated by Irving Fisher's famous 1929 comment about stocks having reached a permanently high plateau.

Haug: In practice is there any way we can hedge options against jump risk without using other options?

> One can get a partial hedge by delta-hedging with the underlying asset, or with futures, but a more exact hedge does indeed require using other options.

Haug: What is the latest on jump risk research?

> I would say there is increasing recognition that not only is there jump risk, but that it varies over time. My 2000 *J.Econometrics* paper shows that models with stochastic jump intensities fit option prices better, and there has been some time series research by, e.g., Eraker, Johannes and Polson supporting that model. We're making substantial progress in devising methods of estimating models with stochastic volatility and/or jump risk directly from time series data – as opposed to calibrating the models from option prices. Finally, the models with Lévy processes are exploring alternate fat-tailed specifications,

and are using random time changes to capture stochastic volatility as well.

Haug: In some recent research, "Hedging the Smirk" you are looking at a "model free" method for inferring deltas and gammas from the market. What is the main idea behind this and how are these deltas and gammas different from the Black-Scholes-Merton ones?

The paper came out of some consulting work I did some time ago: how do you hedge options when the volatility smirk clearly indicates that the Black-Scholes-Merton model is wrong? The paper exploits Euler's theorem: for any homogeneous option pricing function, we can compute deltas (and gammas) from the observable cross-sectional sensitivity of option prices to the strike price. Most of our option pricing models, for stochastic volatility or jump risk, possess this homogeneity property. Translated into implicit volatilities, it creates simple correction terms for deltas and gammas relative to the BSM values for cases when the implicit volatility curve is not perfectly flat.

Haug: The Black-Scholes-Merton delta for deep-out-of-the-money options are extremely sensitive to the volatility used to calculate the delta, do your "model free" method helps here as well? Or do you have any suggestions to how to solve this problem, jump models?

Unfortunately, no; the procedure only identifies how to hedge against delta or gamma risk. It can be used to hedge against pure jump risk; delta-gamma hedging works pretty well for that. However, asset price jumps are typically accompanied by implicit volatility jumps, especially in stock index options, and a delta-hedged position will still have vega risk. Delta-gamma neutral positions will do better; if you neutralize the gamma, you neutralize much of the vega as well.

Haug: When the market is in a normal period; low or moderate volatility individual stocks tend to jump in different directions. So by having a portfolio of stocks some short and some long I will be more or less hedged, some stocks will typically jump up while other jumps down, but in a market crash all stocks tend to jump in the same direction, can you comment on this, how can model and hedge such a behavior?

I gather there is some interest in "dispersion trading" to exploit this: buying portfolios of stock options and writing stock index options. Since individual stock options generally appear to be fairly priced while stock index options are overpriced, it's been a profitable strategy historically. However, it's not a riskless strategy, given that a portfolio of options is not the same as an option on a portfolio. One is betting on the overall level of idiosyncratic stock risk.

Haug: You have also studied crash aversion, in short what is this about?

That's theoretical work in progress aimed at modeling how the option markets function as a market for trading crash risk. Crash aversion reflects individuals' attitudes towards crash risk; while I express it in utility terms, it is fundamentally equivalent to individuals' subjective beliefs as to the frequency of crashes. For trading to take place, people have to differ in some fashion; and differing crash aversion gives them a motive to trade options. Ultimately, I want to introduce frictions, and model the role of option market makers within this market, but I haven't got there yet.

Haug: You have also done some research on the clearinghouse's default exposure during the 1987 crash, did any of the clearinghouses default at that time and what are the main results for your research on this?

None did, but some came close. The work, with Roger Craine, was more aimed at providing additional analytical tools for assessing default exposure. The CME's SPAN system is essentially a VaR approach: setting margins such that the probability of a margin-exceeding move is less than 1%. We were pointing out that clearing houses might also be interested in how much in additional funds would be needed, conditional on margin being exhausted. In the worst-case scenario, the clearing house's responsibility for covering positions means it is essentially short a strangle: a combination of an OTM call and an OTM put. We have methods of evaluating such exposures, from option prices directly or from time series analysis.

Haug: Have we learned from the past, what is the probability we get clearing house defaults in the future in the case of a giant market crash?

I would say yes. Paul Kupiec found that the CME was especially cautious when setting margins on its S&P 500 futures over 1988-92, in the aftermath of the '87 crash. Whether they are still careful, I can't say; it's not something I monitor.

Haug: What about the large number of OTC derivatives that big banks have against other big banks, if one of the big banks defaults could this lead to a chain reaction, more and more banks going under? Do we have control on the derivatives market?

That's the issue of systemic risk – something the Fed is rather worried about. There has also considerable concern about the large positions taken by hedge funds. I don't know if much is being done at the regulatory level. We do have a rapidly growing market in credit derivatives, to hedge against individual defaults.

Haug: A currency is in one way similar to a share on the government of that country, can implied distributions from FX options be used to predict defaults of government obligations and financial crises for a country?

In principle, yes, but it's a bit indirect. Also, FX options are relatively short-maturity. I'd be more inclined to look at fixed-income markets; the country-specific spread over LIBOR on adjustable-rate notes, for instance.

Haug: How good are modern jumps and stochastic volatility models, what is the empirical research telling us?

We continue to make great progress with stochastic volatility models, and now have a better understanding of volatility dynamics. For instance, we now know that volatility has both short-term and long-term swings. Research into jumps has been evolving more slowly, but using intradaily data is providing some new insights. For instance, whereas the jump-diffusion model appears reasonable for daily data, it's less good for intradaily. The 1987 crash did not occur within 5 seconds, for instance; it took all day to fall. A better model is volatility spikes; realized intradaily volatility jumped from its normal level of about 1% daily, to 12% daily on the day of the crash. Aggregated up to daily data (close to close), that looks like a crash; but the intradaily pattern is different.

Haug: With your research background in jumps, stochastic volatility, crash risk and clearinghouse defaults what would you say about VaR and Sharpe ratios as risk management tools for option and other financial instruments?

It's a starting point, and is better than not having any risk management in place. But one has to be careful, given distributions can be far from normal.

Haug: Who was the first to discover fat tails – was it Mandelbrot?

Mandelbrot certainly drew attention to the issue, in the 1960's. Not many use Mandelbrot's stable Paretian specification, however. Infinite-variance processes are rather hard to use in finance, and the prediction that weekly or monthly returns should be as fat-tailed as daily returns appears counterfactual. Merton drew attention to jump-diffusion processes, which are certainly easier to work with. Currently, there's a lot of interest in Lévy processes, which predate Mandelbrot.

Haug: You have spent your career in academia have you ever considered working as an option trader trying to apply your models and ideas in making big bucks and driving sports cars?

I did work for a couple of years at the First National Bank of Chicago, before returning to Princeton for my doctorate, so my career hasn't been purely academic. But yes, it has been mostly academic. And no, I have not considered working as an option trader. I would assume that requires a wholly different set of skills.

Haug: Have you experienced any jumps in stochastic process life integral so far, and if so where they positive or negative.

Moving from Wharton to the University of Iowa was certainly a shock, but overall it has worked out well. I grew up in a college town, in Sewanee, Tennessee, and the life style is comfortable. My family certainly enjoys it.

Haug: Do you have any hobbies outside quantitative finance?

Mostly reading; I'm a great history buff. I used to play tennis, but do less of that, now. I enjoy cycling, and generally try to arrange a cycling trip when I go to Europe.

Haug: In your view where are we in the evolution of quantitative finance?

At this point, I think we're reaching diminishing returns in terms of modeling univariate stochastic processes. We can argue over precise fat-tailed specifications, jumps versus other Lévy, for instance, but the important issue was just recognizing that there are indeed fat tails. Similarly, I'm not sure how much further we can push our stochastic volatility models. There's probably room for improvement in modeling multivariate processes; correlation risk, for instance.

I think we need more work on the finance side, as opposed to the math side. Derivatives permit individuals to trade other risks: credit risk, jump risk and volatility risk, for example. How well do the markets for these risks work? Judging from the profit opportunities from selling OTM puts on stock indexes, not very efficiently.

Those with predominantly a math background feel comfortable with stochastic processes and continuous-time finance; but I see a certain amount of hand-waving when it comes to the transformation from objective to risk-neutral distributions. The use of Esscher transforms, for instance; they work, but what's the rationale? The transformation of probability measure is the issue of compensation for risk – a fundamental finance question. If it's too high, as it appears to be, why isn't money pouring in to exploit the opportunities? I think we need more realistic models of market structure – in particular, of the market makers. But that's an academic's perspective; the practitioners can just go on happily making money off of the market inefficiencies.