Lecture 3: Adding Jumps

Jim Gatheral, Merrill Lynch

Case Studies in Financial Modelling Course Notes,
Courant Institute of Mathematical Sciences,
Fall Term, 2004

7 Jump Diffusion

7.1 Why Jumps are Needed

In section 6, we indicated the possibility that jumps might explain why the skew is so steep for very short expirations and why the very short-dated term structure of skew is inconsistent with any stochastic volatility model. Another indication that jumps might be necessary to explain the volatility surface comes from Table 1. There, we see that there are bids of 0.05 for 750 strike puts and 925 strike calls *with only two days to go*! Given that volatility is around 2% per day according to Figure 1 (2% daily is equivalent to roughly 32% annualized volatility), a 116 point move in the index corresponds to roughly 5 standard deviations. The probability of a normally distributed variable making such a move is about one in a million.

Just as strikingly, in table 2, we see that there is a 5 cent bid for options 45 points out-of-the-money which have almost expired. Recall that the final payoff of SPX options is set at the opening of trading on the following day (September 20 in this case). Historically, about 40% of the variance of SPX is from overnight moves. Then a 45 point move corresponds to around 4.2 standard deviations. The probability of a normally distributed variable making such a move is about one in a hundred thousand. And these 5 cent bids are only bids; one might suppose that actual trades would take place somewhere between the bid and the offer.

In fact, high bids for options that would require an extreme move to end up in-the-money are just another manifestation of the extreme short-end skew in the SPX market just prior to expiration. From the perspective of a trader, the explanation is straightforward: large moves do sometimes occur and it makes economic sense to bid for out-of-the-money options – at the very least to cover existing risk.

It is easy to see why extreme short-end skews are incompatible with stochastic volatility; if the underlying process is a diffusion and volatility of volatility is reasonable, volatility should be near constant on a very short timescale. Then returns should be roughly normally distributed and the skew should be quite flat.

To make this concrete, in Figure 1, we superimpose observed implied volatilities with the implied volatility smile generated by the Heston model with the BCC parameters of Homework 1 (Bakshi, Cao, and Chen (1997)):

$$\eta = 0.39;\ \rho = -0.64;\ \lambda = 1.15$$

Table 1: September 2002 expiration option prices as of September 18, 2002. SPX is trading at 866.

Strike	Call Bid	Call Ask	Put Bid	Put Ask
750	114.00	118.00	0.05	0.25
775	89.00	93.00	0.10	0.45
780	84.00	88.00	0.10	0.45
790	74.10	78.10	0.15	0.45
800	64.20	68.20	0.15	0.40
810	54.40	58.40	-	1.00
820	44.70	48.70	0.25	1.25
825	39.90	43.90	0.45	1.45
830	35.20	39.20	0.75	1.75
840	26.20	30.20	1.75	2.30
850	18.00	21.50	3.00	4.60
860	10.90	13.00	5.50	7.40
870	5.90	8.90	10.00	12.90
875	4.50	4.80	12.50	16.00
880	2.95	4.10	15.90	19.40
885	1.95	2.95	19.80	23.30
890	1.10	1.50	23.60	27.60
900	0.40	0.80	32.60	36.60
905	0.05	1.00	37.40	41.40
910	-	0.50	42.20	46.20
915	-	0.30	47.10	51.10
920	0.05	0.50	52.10	56.10
925	0.05	0.25	57.10	61.10

Table 2: September 2002 expiration option prices as of Thursday September 19, 2002 at 4PM. SPX is trading at 843.

Strike	**Call Bid**	**Call Ask**	**Put Bid**	**Put Ask**
800	41.20	45.20	0.05	0.20
810	31.50	35.50	0.15	0.30
820	22.10	26.10	0.65	1.00
825	18.00	21.20	1.00	1.90
830	13.80	17.00	1.95	2.95
840	7.00	9.00	4.30	5.90
850	2.00	2.35	9.30	10.00
860	0.60	0.65	16.10	17.80
870	-	0.40	25.10	29.10
875	0.10	0.20	30.10	34.10
880	-	0.50	35.10	39.10
885	-	0.30	40.00	44.00
890	0.05	0.10	46.00	49.00

Figure 1: Graph of SPX implied volatilities on September 18, 2002. SPX is trading at 866. Red points are offers and blue points are bids. The green line is a non-linear fit to the data. The red line represents the Heston skew with BCC parameters.

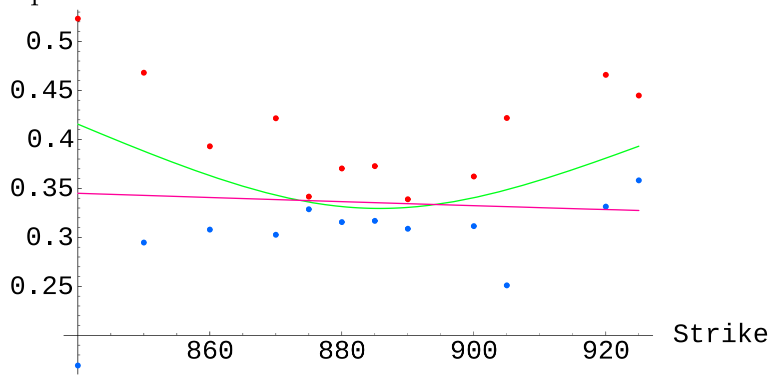

7.2 Derivation of the Valuation Equation

As in Wilmott (1998), we assume the stock price follows the SDE

$$dS = \mu S\,dt + \sigma S\,dZ + (J-1)S\,dq \tag{43}$$

where the Poisson process

$$dq = \begin{cases} 0 & \text{with probability } 1-\lambda(t)\,dt \\ 1 & \text{with probability } \lambda(t)\,dt \end{cases}$$

When $dq = 1$, the process jumps from S to JS. We assume that the Poisson process dq and the Brownian motion dZ are independent.

As in the stochastic volatility case, we derive a valuation equation by considering the hedging of a contingent claim. We make the (unrealistic) assumption at this stage that the jump size J is known in advance.

Whereas in the stochastic volatility case, the second risk factor to be hedged was the random volatility, in this case, the second factor is the jump. So once again, we set up a portfolio Π containing the option being priced whose value we denote by $V(S,v,t)$, a quantity $-\Delta$ of the stock and a quantity $-\Delta_1$ of another asset whose value V_1 also depends on the jump.

We have

$$\Pi = V - \Delta\,S - \Delta_1\,V_1$$

The change in this portfolio in the time interval dt is given by

$$\begin{aligned} d\Pi &= \left\{\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2\,S^2\frac{\partial^2 V}{\partial S^2}\right\} dt - \Delta_1\left\{\frac{\partial V_1}{\partial t} + \frac{1}{2}\sigma^2\,S^2\frac{\partial^2 V_1}{\partial S^2}\right\} dt \\ &\quad + \left\{\frac{\partial V}{\partial S} - \Delta_1\frac{\partial V_1}{\partial S} - \Delta\right\} dS^c \\ &\quad + \{V(JS,t) - V(S,t) - \Delta_1(V_1(JS,t) - V_1(S,t)) - \Delta(J-1)S\}\ dq \end{aligned}$$

where $S^c(t)$ is the continuous part of $S(t)$ (adding back all the jumps that occurred up to time t).

To make the portfolio instantaneously risk-free, we must choose

$$\frac{\partial V}{\partial S} - \Delta_1\frac{\partial V_1}{\partial S} - \Delta = 0$$

to eliminate dS terms, and

$$V(JS,t) - V(S,t) - \Delta_1(V_1(JS,t) - V_1(S,t)) - \Delta(J-1)S = 0$$

to eliminate dq terms. This leaves us with

$$\begin{aligned}
d\Pi &= \left\{\frac{\partial V}{\partial t}+\frac{1}{2}\sigma^2\,S^2\frac{\partial^2 V}{\partial S^2}\right\}dt-\Delta_1\left\{\frac{\partial V_1}{\partial t}+\frac{1}{2}\sigma^2\,S^2\frac{\partial^2 V_1}{\partial S^2}\right\}dt\\
&= r\,\Pi\,dt\\
&= r(V-\Delta S-\Delta_1 V_1)\,dt
\end{aligned}$$

where we have used the fact that the return on a risk-free portfolio must equal the risk-free rate r which we will assume to be deterministic for our purposes. Collecting all V terms on the left-hand side and all V_1 terms on the right-hand side, we get

$$\frac{\frac{\partial V}{\partial t}+\frac{1}{2}\sigma^2\,S^2\frac{\partial^2 V}{\partial S^2}+rS\frac{\partial V}{\partial S}-rV}{\delta V-(J-1)S\frac{\partial V}{\partial S}}=\frac{\frac{\partial V_1}{\partial t}+\frac{1}{2}\sigma^2\,S^2\frac{\partial^2 V_1}{\partial S^2}+rS\frac{\partial V_1}{\partial S}-rV_1}{\delta V_1-(J-1)S\frac{\partial V_1}{\partial S}}$$

where we have defined $\delta V \equiv V(JS,t)-V(S,t)$.

Continuing exactly as in the stochastic volatility case, the left-hand side is a function of V only and the right-hand side is a function of V_1 only. The only way that this can be is for both sides to be equal to some function of the independent variables S and t which we will suggestively denote by $-\lambda$. We deduce that

$$\begin{aligned}
\frac{\partial V}{\partial t} &+ \frac{1}{2}\sigma^2\,S^2\frac{\partial^2 V}{\partial S^2}+rS\frac{\partial V}{\partial S}-rV\\
&+ \lambda(S,t)\left\{V(JS,t)-V(S,t)-(J-1)S\frac{\partial V}{\partial S}\right\}=0
\end{aligned}\tag{44}$$

To interpret $\lambda(S,t)$, consider the value P of an asset that pays \$1 at time T if there is no jump and zero otherwise. Our assumption that the jump process is independent of the stock price process implies that

$$\frac{\partial P}{\partial S}=0$$

Also, we must have $P(JS,t)=0$. Substituting into equation (44) gives

$$\frac{\partial P}{\partial t}-rP-\lambda(S,t)\,P=0$$

Since (by assumption) P is independent of S, so must λ be and the solution is $P(t)=\exp\left\{-\int_t^T(r+\lambda(t'))\,dt'\right\}$. We immediately recognize $\lambda(t)$ as the hazard rate of the Poisson process (the *pseudo-probability* per unit time that a jump occurs). We emphasize pseudo-probability because this is in no sense the actual probability (whatever that means) that a jump will occur: it is the value today of a financial asset.

Uncertain jump size

To derive equation (44), we assumed that we knew in advance what the jump size would be. Of course this is neither realistic nor practical. Jump-diffusion models typically specify a distribution of jump sizes. How would this change equation (44)?

It is easy to see that adding another jump with a different size would require one more hedging asset in the replication argument. Allowing the jump size to be any real number with some distribution would require an infinite number of hedging assets. We see that in this case, the replication argument falls apart: such jump-diffusion models have no replicating hedge.

This is the major drawback of jump-diffusion models: there is no replicating portfolio and so there is no self-financing hedge even in the limit of continuous trading. However, looking on the bright side, if we believe in jumps (as we must given the empirical evidence), options are no longer redundant assets which may be replicated using stocks and bonds and by extension, option traders can be seen to have genuine social value.

To extend equation (44) to the case of jumps of uncertain size, we assume that the risk-neutral process is still jump-diffusion with jumps independent of the stock price. Under the risk-neutral measure, the expected return of any asset is the risk-free rate. Taking expectations of equation (43), we find that

$$\mathbb{E}[dS] = r\,S\,dt = \mu\,S\,dt + \mathbb{E}[J-1]\,S\,\lambda(t)\,dt$$

It follows that the risk-neutral drift is given by $\mu = r + \mu_J$ with

$$\mu_J = -\lambda(t)\,\mathbb{E}[J-1]$$

Just as in the derivation of the Black-Scholes equation, we must have $\mathbb{E}[dV] = r\,V\,dt$. Applying Itô's Lemma and taking expectations under the risk-neutral measure gives

$$
\begin{aligned}
\mathbb{E}[dV] &= r\,V\,dt \\
&= \left\{\frac{\partial V}{\partial t} + r\,S\,\frac{\partial V}{\partial S} + \frac{1}{2}\sigma^2\,S^2\frac{\partial^2 V}{\partial S^2}\right\} dt \\
&\quad + \lambda(t)\,\mathbb{E}[V(JS,t) - V(S)]\,dt + \mu_J\,S\,\frac{\partial V}{\partial S}\,dt
\end{aligned}
$$

Rearranging, we obtain the following equation for valuing financial assets under jump-diffusion:

$$
\begin{aligned}
\frac{\partial V}{\partial t} &+ \frac{1}{2}\sigma^2\,S^2\frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV \\
&+ \lambda(t)\left\{\mathbb{E}\left[V(JS,t) - V(S,t)\right] - \mathbb{E}\left[J-1\right]S\frac{\partial V}{\partial S}\right\} = 0
\end{aligned}
\tag{45}
$$

Once again for emphasis, the expectations in equation (45) are under the risk neutral measure. In order to value derivative assets, we concern ourselves only with the values that the market assigns to claims that pay in the event of a jump; actual probabilities don't enter at all.

8 Characteristic Function Methods

Unlike the partial differential equations (PDEs) we are used to solving in derivatives valuation problems, equation (45) is an example of an partial integro-differential equation (PIDE). The integration over all possible jump-sizes introduces non-locality. Such equations can be solved using extensions of numerical PDE techniques but the most natural approach is to use Fourier transform (characteristic function) methods.

First, we review Lévy processes.

8.1 Lévy Processes

With constant hazard rate λ, the logarithmic version of the jump-diffusion process (43) for the underlying asset is an example of a Lévy Process.

Definition. *A Lévy process is a continuous in probability, cadlag stochastic process $x(t)$, $t > 0$ with independent and stationary increments and $x(0) = 0$.*

It turns out that any Lévy process can be expressed as the sum of a linear drift term, a Brownian motion and a jump process. This plus the independent increment property leads directly to the following representation for the characteristic function.

The Lévy-Khintchine Representation

If x_t is a Lévy process, and if the Lévy density $\mu(\xi)$ is suitably well-behaved at the origin, its characteristic function $\phi_T(u) \equiv \mathbb{E}\left[e^{iux_T}\right]$ has the representation

$$\phi_T(u) = \exp\left\{iu\omega T - \frac{1}{2}u^2\,\sigma^2 T + T\int\left[e^{iu\xi} - 1\right]\mu(\xi)\,d\xi\right\} \tag{46}$$

To get the drift parameter ω, we impose that the risk-neutral expectation of the stock price be the forward price. With our current assumption of zero interest rates and dividends, this translates to imposing that

$$\phi_T(-i) = \mathbb{E}\left[e^{x_T}\right] = 1$$

Here, $\int \mu(\xi)\,d\xi = \lambda$, the Poisson intensity or mean jump arrival rate also known as the *hazard* rate.

8.2 Examples of characteristic functions for specific processes

Before proceeding to solve equation (45) for a particular specification of the jump process, we exhibit some examples of characteristic functions for processes with which we are already familiar.

Example 1: Black-Scholes

The characteristic function for a exponential Brownian motion with volatility σ is given by

$$\phi_T(u) = \mathbb{E}\left[e^{iux_T}\right] = \exp\left\{-\frac{1}{2}u(u+i)\sigma^2 T\right\}$$

We can get this result by performing the integration explicitly or directly from the Lévy-Khintchine representation.

Example 2: Heston

The Heston process is very path-dependent; increments are far from independent and it is not a Lévy process. However, we have already computed its characteristic function. From Section 3.2, we see that the characteristic function of the Heston process is given by

$$\phi_T(u) = \exp\left\{C(u,T)\,\bar{v} + D(u,T)\,v\right\}$$

with $C(u,T)$ and $D(u,T)$ as defined there.

Example 3: Merton's Jump-Diffusion Model

Finally, this is the case we are really interested in. The jump-size J is assumed to be lognormally distributed with mean log-jump α and standard deviation δ so that the stock price follows the SDE

$$dS = \mu S\,dt + \sigma S\,dZ + (e^{\alpha+\delta\epsilon} - 1)S\,dq$$

with $\epsilon \sim N(0,1)$. Then

$$\mu(\xi) = \frac{\lambda}{\sqrt{2\pi\delta^2}}\exp\left\{-\frac{(\xi-\alpha)^2}{2\delta^2}\right\}$$

By applying the Lévy-Khintchine representation (46), we see that the characteristic function is given by

$$\begin{aligned}
\phi_T(u) &= \exp\left\{iu\omega T - \frac{1}{2}u^2\,\sigma^2 T + T\int\left[e^{iu\xi} - 1\right]\frac{\lambda}{\sqrt{2\pi\delta^2}}\exp\left\{-\frac{(\xi-\alpha)^2}{2\delta^2}\right\}\,d\xi\right\} \\
&= \exp\left\{iu\omega T - \frac{1}{2}u^2\,\sigma^2 T + \lambda T\left(e^{iu\alpha - u^2\delta^2/2} - 1\right)\right\} \qquad (47)
\end{aligned}$$

To get ω, we impose $\phi_T(-i) = 1$ so that

$$\exp\left\{\omega T + \frac{1}{2}\,\sigma^2 T + \lambda T\left(e^{\alpha+\delta^2/2} - 1\right)\right\} = 1$$

which gives

$$\omega = -\frac{1}{2}\,\sigma^2 - \lambda\left(e^{\alpha+\delta^2/2} - 1\right)$$

Unsurprisingly, we get the lognormal case back when we set $\alpha = \delta = 0$.

Alternatively, we can get the characteristic function for jump-diffusion directly by substituting $\phi_T(u) = e^{\psi(u)\,T}$ into equation (45). With $y \sim N(\alpha, \delta)$, we obtain

$$\begin{aligned}\psi(u) &= -\frac{1}{2}u(u+i)\sigma^2 - \lambda\left\{\mathbb{E}\left[e^{iuy}-1\right] + iu\,\mathbb{E}\left[e^y - 1\right]\right\} \\ &= -\frac{1}{2}u(u+i)\sigma^2 - \lambda\left\{\left(e^{iu\alpha - u^2\,\delta^2/2} - 1\right) + iu\,\left(e^{\alpha+\delta^2/2} - 1\right)\right\}\end{aligned}$$

which gives an expression for $\phi_T(u)$ identical to the one already derived in equation (47).

8.3 Computing Option Prices From the Characteristic Function

It turns out (see Carr and Madan (1999) and Lewis (2000)) that it is quite straightforward to get option prices by inverting the characteristic function of a given stochastic process (if it is known in closed-form).

The formula we will use is a special case of formula (2.10) of Lewis (as usual we assume zero interest rates and dividends):

$$C(S,K,T) = S - \sqrt{SK}\frac{1}{\pi}\int_0^\infty \frac{du}{u^2+\frac{1}{4}}\,\mathbf{Re}\left[e^{-iuk}\phi_T\left(u - i/2\right)\right] \tag{48}$$

with $k = \ln\left(\frac{K}{S}\right)$. A proof of this formula is given in Appendix A.

8.4 Computing Implied Volatility

Equation (48) allows us to derive an elegant implicit expression for the Black-Scholes implied volatility of an option in any model for which the characteristic function is known.

Substituting the characteristic function for the Black-Scholes process into (48) gives

$$C_{BS}(S,K,T) = S - \sqrt{SK}\frac{1}{\pi}\int_0^\infty \frac{du}{u^2+\frac{1}{4}}\,\mathbf{Re}\left[e^{-iuk}e^{-\frac{1}{2}\left(u^2+\frac{1}{4}\right)\sigma_{BS}^2 T}\right]$$

Then, from the definition of implied volatility, we must have

$$\int_0^\infty \frac{du}{u^2+\frac{1}{4}}\,\mathbf{Re}\left[e^{-iuk}\left(\phi_T\left(u-i/2\right) - e^{-\frac{1}{2}\left(u^2+\frac{1}{4}\right)\sigma_{BS}^2 T}\right)\right] = 0 \tag{49}$$

Equation (49) gives us a simple but implicit relationship between the implied volatility surface and the characteristic function of the underlying stock process. In particular, we may efficiently compute the structure of at-the-money implied volatility and the at-the-money volatility skew in terms of the characteristic function (at least numerically) without having to explicitly compute option prices.

8.5 Computing the At-the-money Volatility Skew

Assume ϕ_T does not depend on spot S and hence not on k. (This is the case in all examples we have in mind.) Then differentiating (49) with respect to k and evaluating at $k = 0$ gives

$$\int_0^\infty du \left\{ \frac{u\,\mathbf{Im}\left[\phi_T(u - i/2)\right]}{u^2 + \frac{1}{4}} + \frac{1}{2}\left.\frac{\partial w_{BS}}{\partial k}\right|_{k=0} e^{-\frac{1}{2}\left(u^2+\frac{1}{4}\right)w_{BS}(0,T)} \right\} = 0$$

Then, integrating the second term explicitly we get

$$\left.\frac{\partial \sigma_{BS}}{\partial k}\right|_{k=0} = -e^{\frac{\sigma_{BS}^2 T}{8}} \sqrt{\frac{2}{\pi}} \frac{1}{\sqrt{T}} \int_0^\infty du\, \frac{u\,\mathbf{Im}\left[\phi_T(u - i/2)\right]}{u^2 + \frac{1}{4}} \tag{50}$$

Example 1: Black-Scholes

$$\mathbf{Im}\left[\phi_T(u - i/2)\right] = \mathbf{Im}\left[e^{-\frac{1}{2}\left(u^2+1/4\right)\sigma^2 T}\right] = 0$$

Then, in the Black-Scholes case,

$$\left.\frac{\partial \sigma_{BS}(k,T)}{\partial k}\right|_{k=0} = 0 \;\; \forall T > 0$$

Example 2: Merton's Jump-Diffusion Model (JD)

Write

$$\phi_T(u) = e^{-\frac{1}{2}u(u+i)\sigma^2 T}\, e^{\psi(u)T}$$

with $\psi(u) = -\lambda i u \left(e^{\alpha+\delta^2/2} - 1\right) + \lambda\left(e^{iu\alpha - u^2\delta^2/2} - 1\right)$

Then

$$\mathbf{Im}\left[\phi_T(u - i/2)\right] = e^{-\frac{1}{2}\left(u^2+\frac{1}{4}\right)\sigma^2 T}\,\mathbf{Im}\left[e^{\psi(u-i/2)T}\right]$$

8.6 How jumps impact the volatility skew

By substituting the jump-diffusion characteristic function (47) into our expressions (49) and (50) for the implied volatility and ATM volatility skew respectively, we can investigate the impact of jumps on the volatility surface for various numerical choices of the parameters.

Figure 2: The 3 month volatility smile for various choices of jump-diffusion parameters.

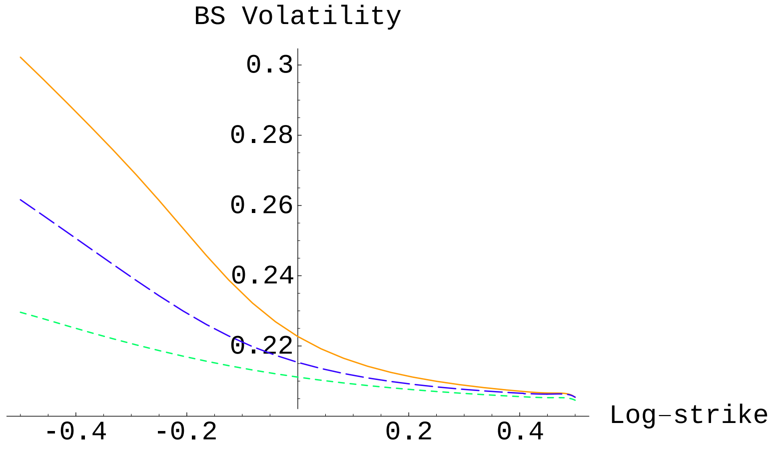

Figure 3: The term structure of ATM variance skew for various choices of jump-diffusion parameters.

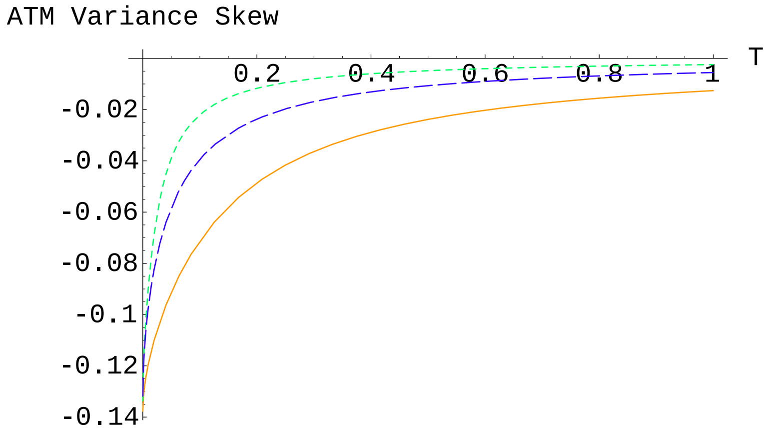

Table 3: Parameters used to generate Figures 2 and 3.

Color	σ	λ	α	δ
Solid orange	0.2	0.5	-0.15	0.05
Dashed green	0.2	1.0	-0.07	0.00
Long-dashed blue	0.2	1.0	-0.07	0.05

8.6.1 The Decay of Skew Due to Jumps

We can see from Figure 2 that the volatility skew decays very rapidly in a jump-diffusion model beyond a certain time to expiration. To estimate this characteristic time, we note that prices of European options depend only on the final distribution of stock prices and if the jump size is of the order of only one standard deviation $\sigma\sqrt{T}$, a single jump has little impact on the shape of this distribution. If there are many small jumps, returns will be hard to distinguish from normal over a reasonable time interval. We compute the characteristic time T^* by equating

$$-\left(e^{\alpha+\delta^2/2}-1\right)\approx\sigma\sqrt{T^*}$$

8.6.2 Skew behavior under jump-diffusion as $T \to 0$

Consider the value of an option under jump-diffusion with a short time ΔT to expiration. Because the time to expiration is very short, the probability of having more than one jump is negligible. Because the jump is independent of the diffusion, the value of the option is just a superposition of the value conditional on the jump and the value conditional on no jump. Without loss of generality, suppose the stock price jumps down from S to JS when the jump occurs. Then

$$\begin{aligned} C_J(S,K,\Delta T) &\approx (1-\lambda\Delta T)\,C_{BS}(Se^{\mu_J\Delta T},K,\Delta T)+\lambda\Delta T\,C(JS,K,\Delta T) \\ &= C_{BS}(Se^{\mu_J\Delta T},K,\Delta T)+O(\Delta T) \end{aligned} \tag{51}$$

where J is the size of the jump, $C_J(.)$ represents the value of the option under jump diffusion and $\mu_J=-\lambda(e^{\alpha+\delta^2/2}-1)$ is the adjustment to the risk-neutral drift for jumps. Here, we neglected the second term in equation (51) by assuming that the mean jump is downwards and the probability of the option being in-the-money is negligibly small after the jump.

We want to compute the at-the-money variance skew

$$\left.\frac{\partial \sigma_{BS}^2}{\partial k}\right|_{k=0}$$

To do this note that

$$\frac{\partial C_J}{\partial k} = \frac{\partial C_{BS}}{\partial k} + \frac{\partial C_{BS}}{\partial \sigma_{BS}}\frac{\partial \sigma_{BS}}{\partial k}$$

so

$$\left.\frac{\partial \sigma_{BS}}{\partial k}\right|_{k=0} = \left[\frac{\partial C_J}{\partial k} - \frac{\partial C_{BS}}{\partial k}\right]\left.\left(\frac{\partial C_{BS}}{\sigma_{BS}}\right)^{-1}\right|_{k=0}$$

Now, for an at-the-money option,

$$\left.\frac{\partial C_{BS}}{\sigma_{BS}}\right|_{k=0} \approx \frac{S}{\sqrt{2\pi}}\sqrt{\Delta T}$$

and from equation (51)

$$\begin{aligned}
\left.\frac{1}{S}\left[\frac{\partial C_J}{\partial k} - \frac{\partial C_{BS}}{\partial k}\right]\right|_{k=0} &\approx -N\left(+\frac{\mu_J \Delta T}{\sigma\sqrt{\Delta T}} - \frac{1}{2}\sigma\sqrt{\Delta T}\right) + N\left(-\frac{1}{2}\sigma\sqrt{\Delta T}\right) \\
&\approx -\frac{1}{\sqrt{2\pi}}\frac{\mu_J}{\sigma}\sqrt{\Delta T}
\end{aligned}$$

Then, for small ΔT,

$$\left.\frac{\partial \sigma_{BS}^2}{\partial k}\right|_{k=0} \approx -2\,\mu_J \tag{52}$$

We see that in a jump-diffusion model, if the mean jump-size is sufficiently large relative to its standard deviation, the at-the-money variance skew is given directly by twice the jump compensator μ_J.

To see how well these approximate computations explain Figures 2 and 3, the characteristic time T^* and the time zero skew ψ_0 for each choice of parameters are presented in Table 4.

Summarizing the results, we note that the jump compensator (or expected move in the stock price due to jumps) drives the skew in the short-expiration limit while the decay of ATM skew is driven by the expected jump size.

Table 4: Interpreting Figures 2 and 3.

Color	σ	λ	α	δ	$\mathbf{T}^*$	$\psi_\mathbf{0}$
Solid orange	0.2	0.5	-0.15	0.05	0.69	-0.133
Dashed green	0.2	1.0	-0.07	0.00	0.34	-0.135
Long-dashed blue	0.2	1.0	-0.07	0.05	0.33	-0.133

9 Stochastic Volatility plus Jumps

9.1 Stochastic Volatility plus Jumps in the Underlying Only (SVJ)

Since jumps generate a steep short-dated skew that dies quickly with time to expiration and stochastic volatility models don't generate enough skew for very short expirations but more or less fit for longer expirations (see Lecture 2), it is natural to try to combine stock price jumps and stochastic volatility in one model.

Suppose we add a simple Merton-style lognormally distributed jump process to the Heston process. By substitution into the valuation equation, it is easy to see that the characteristic function for this process is just the product of Heston and jump characteristic functions. Denoting the jump intensity (or hazard rate) by λ_J, we obtain

$$\phi_T(u) = e^{C(u,T)\,\bar{v}+D(u,T)\,v}\; e^{\psi(u)T}$$

with $\psi(u) = -\lambda_J iu\left(e^{\alpha+\delta^2/2}-1\right)+\lambda_J\left(e^{iu\alpha-u^2\delta^2/2}-1\right)$ and $C(u,T)$, $D(u,T)$ are as before.

Again, we may substitute this functional form into equations (49) and (50) to get the implied volatilities and at-the-money volatility skew respectively for any given expiration.

Figure 4 plots the at-the-money variance skew corresponding to the Bakshi-Cao-Chen SVJ model fit together with the sum of the Heston and jump-diffusion at-the-money variance skews with the same parameters (see Table 5). We see that (at least with this choice of parameters), not only does the characteristic function factorize but the at-the-money variance skew is almost additive. One practical consequence of this is that the Heston parameters can be fitted fairly robustly using longer dated options and then jump parameters can be found to generate the required extra skew for short-dated options. Figure 5 plots the at-the-money variance skew corresponding to the SVJ model vs the Heston model skew for short-dated options, highlighting the small difference.

Figure 4: The green line is a graph of the at-the-money variance skew in the SVJ model with BCC parameters vs time to expiration. The dashed blue line represents the sum of at-the-money Heston and jump-diffusion skews with the same parameters.

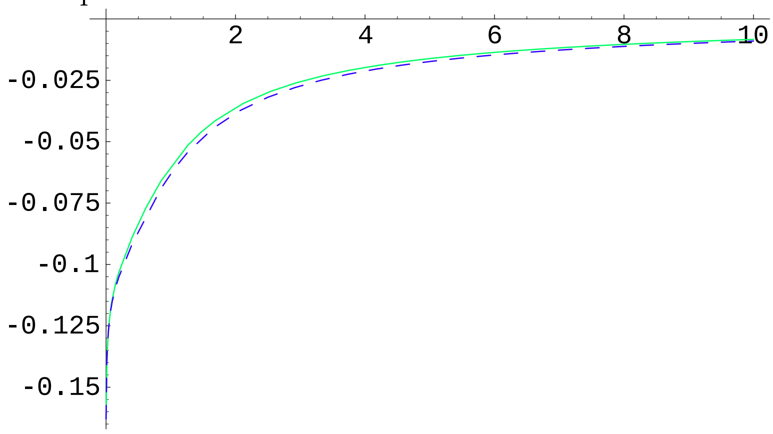

Figure 5: The green line is a graph of the at-the-money variance skew in the SVJ model with BCC parameters vs time to expiration. The dashed red line represents the at-the-money Heston skew with the same parameters.

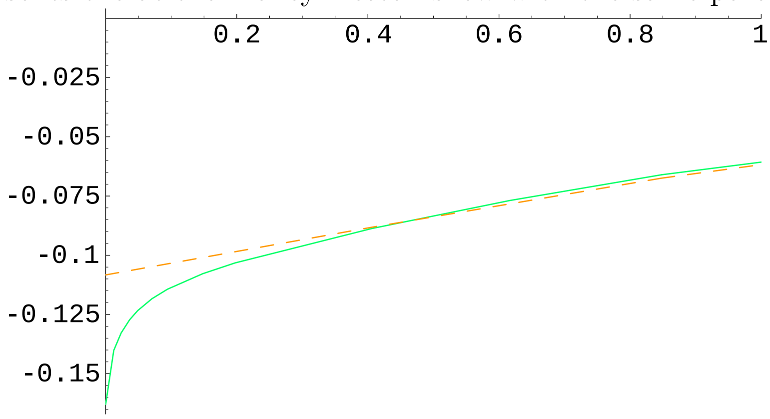

However in the SVJ model, after the stock price has jumped, the volatility will stay unchanged because the jump process is uncorrelated with the volatility process. This is inconsistent with both intuition and empirically observed properties of the time series of asset returns; in practice, after a large move in the underlying, implied volatilities always increase substantially (*i.e.* they jump).

9.2 Some Empirical Fits to the SPX Volatility Surface

There are only 4 parameters in the jump-diffusion model: the volatility σ, λ_J, α and δ so it's not in principle difficult to perform a fit to option price data. The SVJ model obviously fits the data better because it has more parameters and it's not technically that much harder to perform the fit.

Various authors (for example Andersen and Andreasen (2000) and Duffie, Pan, and Singleton (2000)) have fitted JD and SVJ models to SPX data. Their results are summarized in Table 5.

Table 5: Various fits of jump-diffusion style models to SPX data. JD means Jump Diffusion and SVJ means Stochastic Volatility plus Jumps.

Author(s)	Model	λ	η	ρ	$\bar{v}$	$\lambda_{\mathbf{J}}$	α	δ
AA	JD	NA	NA	NA	0.032	0.089	-0.8898	0.4505
BCC	SVJ	2.03	0.38	-0.57	0.04	0.61	-0.09	0.14
M	SVJ	1.0	0.8	-0.7	0.04	0.5	-0.15	0
DPS	SVJ	3.99	0.27	-0.79	0.014	0.11	-0.12	0.15

Author(s)	Reference	Data from
AA	Andersen and Andreasen (2000)	April 1999
BCC	Bakshi, Cao, and Chen (1997)	June 1988 – May 1991
M	Matytsin (1999)	1999
DPS	Duffie, Pan, and Singleton (2000)	November 1993

Note first that these estimates all relate to different dates so in principle, we can't expect the volatility surfaces they generate to be the same shape. Nevertheless, the shape of the SPX volatility surface doesn't really change much over time so it does make some sense to compare them.

9.3 Stochastic volatility with Simultaneous Jumps in Stock Price and Volatility (SVJJ)

As we noted earlier in Section 9.1, it is unrealistic to suppose that the instantaneous volatility wouldn't jump if the stock price were to jump. Conversely, adding a simultaneous upward jump in volatility to jumps in the stock price allows us to maintain the clustering property of stochastic volatility models: recall that "large moves follow large moves and small moves follow small moves".

In Matytsin (1999) and Matytsin (2000), Andrew Matytsin describes a model that is effectively SVJ with a jump in volatility: jumps in the stock price are accompanied by a jump $v \mapsto v + \gamma_v$ in the instantaneous volatility. In that case, the characteristic function is

$$\phi_T(u) = \exp\left\{\hat{C}(u,T)\,\bar{v} + \hat{D}(u,T)\,v\right\} \tag{53}$$

with $C(u,T)$ and $D(u,T)$ given by

$$\begin{aligned}
\hat{C}(u,T) &= C(u,T) + \lambda_J\,T\,\left[e^{iu\alpha - u^2\delta^2/2} I(u,T) - 1 - i\,u\,\left(e^{\alpha+\delta^2/2} - 1\right)\right] \\
\hat{D}(u,T) &= D(u,T)
\end{aligned}$$

where

$$\begin{aligned}
I(u,T) &= \frac{1}{T}\int_0^T e^{\gamma_v D(u,T)}dt \\
&= -\frac{2\gamma_v}{p_+ p_-}\int_0^{-\gamma_v D(u,T)} \frac{e^{-z}dz}{(1+z/p_+)(1+z/p_-)}
\end{aligned}$$

and

$$p_\pm = \frac{\gamma_v}{\eta^2}(\beta - \rho\eta u i \pm d)$$

In the limit $\gamma_v \to 0$, we have $I(u,T) \to 1$ and by inspection, we retrieve the SVJ model. Also, in the limit $T \to 0$, $I(u,T) \to 1$ and in that limit, the SVJJ characteristic function is identical to the SVJ characteristic function. Alternatively, following the heuristic argument of section 8.6.2, the short-dated volatility skew is a function of the jump compensator only and this compensator is identical in the SVJ and SVJJ cases. Intuitively, when the stock price jumps, the volatility jumps but this has no effect in the $T \to 0$

limit because by assumption, an at-the-money option is always out-of-the-money after the jump and its time is zero no matter what the volatility is.

On the other hand, in the $T \to \infty$ limit, the skew should increase because the effective volatility of volatility increases due to (random) jumps in volatility.

By substituting the SVJJ characteristic function (53) into equation (50) for the implied volatility skew with the BCC parameters plus a variance jump of $\gamma_v = 0.1$, we obtain the graphs shown in Figures 6 and 7. We note that the term structure of volatility skew is in accordance with our intuition. In particular, adding a jump in volatility doesn't help explain extreme short-dated volatility skews. However relative to stochastic volatility and SVJ models, it does reduce the volatility of volatility required to fit longer-dated volatility skews even if that comes at the expense of a seemingly even more unreasonable estimate for the average stock price jump.

Figure 6: The orange line is a graph of the at-the-money variance skew in the SVJJ model with BCC parameters vs time to expiration. The short-dashed blue and long-dashed green lines are SVJ and Heston skew graphs respectively with the same parameters.

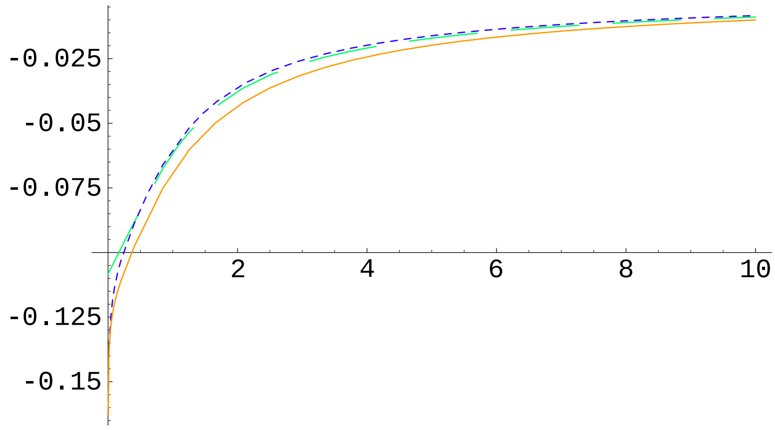

Figure 7: This graph is a short-expiration detailed view of the graph shown in Figure 6.

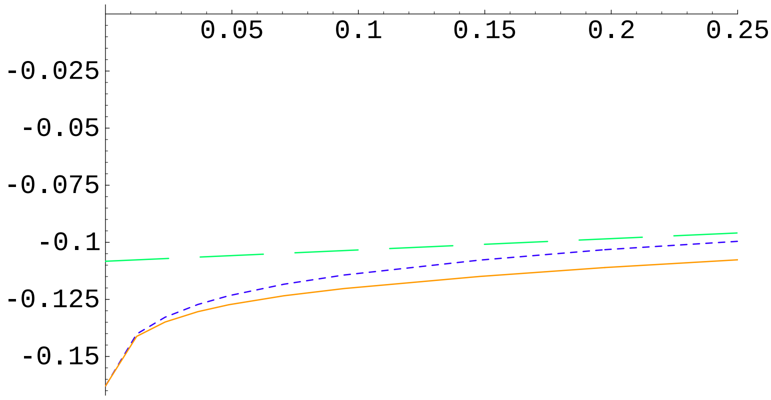

References

Andersen, Leif, and Jesper Andreasen, 2000, Jump-diffusion processes: Volatility smile fitting and numerical methods for option pricing, *Review of Derivatives Research* 4, 231–262.

Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 1997, Empirical performance of alternative option pricing models, *The Journal of Finance* 52, 2003–2049.

Carr, Peter, and Dilip B. Madan, 1999, Option pricing using the fast fourier transform, *The Journal of Computational Finance* 2, 61–73.

Duffie, Darrell, Jun Pan, and Kenneth Singleton, 2000, Transform analysis and asset pricing for affine jump diffusions, *Econometrica* 68, 1343–1376.

Lewis, Alan L., 2000, *Option Valuation under Stochastic Volatility with Mathematica Code* . chap. 2, pp. 34–75 (Finance Press: Newport Beach, CA).

Matytsin, Andrew, 1999, Modelling volatility and volatility derivatives, Columbia Practitioners Conference on the Mathematics of Finance.

——— , 2000, Perturbative analysis of volatility smiles, Columbia Practitioners Conference on the Mathematics of Finance.

Merton, Robert C., 1976, Option pricing when underlying stock returns are discontinuous, *Journal of Financial Economics* 3, 125–144.

Wilmott, Paul, 1998, *Derivatives. The Theory and Practice of Financial Engineering* . chap. 23, pp. 299–304 (John Wiley & Sons: Chichester).

A Proof of Equation (48)

A covered call position has the payoff $\min[S_T, K]$ where S_T is the stock price at time T and K is the strike price of the call. Consider the Fourier transform of this covered call position $G(k,\tau)$ with respect to the log-strike $k \equiv \log(K/F)$ defined by

$$\hat{G}(u,\tau) = \int_{-\infty}^{\infty} e^{iuk} G(k,\tau)\, dx$$

Denoting the current time by t and expiration by T, and setting interest rates and dividends to zero as usual, we have that

$$\begin{aligned}
\frac{1}{S}\hat{G}(u, T-t) &= \int_{-\infty}^{\infty} e^{iuk} \mathbb{E}\left[\min[e^{x_T}, e^k)^+] | x_t = 0\right] dk \\
&= \mathbb{E}\left[\int_{-\infty}^{\infty} e^{iuk} \min[e^{x_T}, e^k)^+]\, dk \,\middle|\, x_t = 0\right] \\
&= \mathbb{E}\left[\int_{-\infty}^{x_T} e^{iuk} e^k\, dk + \int_{x_T}^{\infty} e^{iuk} e^{x_T}\, dk \,\middle|\, x_t = 0\right] \\
&= \mathbb{E}\left[\frac{e^{(1+iu)x_T}}{1+iu} - \frac{e^{(1+iu)x_T}}{iu} \,\middle|\, x_t = 0\right] \text{ only if } 0 < \mathrm{Im}[u] < 1! \\
&= \frac{1}{u(u-i)} \mathbb{E}\left[e^{(1+iu)x_T} \middle| x_t = 0\right] \\
&= \frac{1}{u(u-i)}\, \phi_T(u-i)
\end{aligned}$$

by definition of the characteristic function $\phi_T(u)$. Note that the transform of the covered call value exists only if $0 < \mathrm{Im}[u] < 1$. It is easy to see that this derivation would go through pretty much as above with other payoffs though it is key to note that the region where the transform exists depends on the payoff.

To get the call price in terms of the characteristic function, we express it in terms of the covered call and invert the Fourier transform, integrating along the line $\mathrm{Im}[u] = 1/2$[1]. Then

$$\begin{aligned}
C(S,K,T) &= S - S\frac{1}{2\pi}\int_{-\infty+i/2}^{\infty+i/2} \frac{du}{u(u-i)}\,\phi_T(u-i)\,e^{-iku} \\
&= S - S\frac{1}{2\pi}\int_{-\infty}^{\infty} \frac{du}{(u+i/2)(u-i/2)}\,\phi_T(u-i/2)\,e^{-ik(u+i/2)} \\
&= S - \sqrt{SK}\frac{1}{\pi}\int_{0}^{\infty} \frac{du}{u^2+\frac{1}{4}}\,\mathbf{Re}\left[e^{-iuk}\phi_T\left(u-i/2\right)\right]
\end{aligned}$$

with $k = \ln\left(\frac{K}{S}\right)$.

[1]That's why we chose to express the call in terms of the covered call whose transform exists in this region. Alternatively, we could have used the transform of the call price and Cauchy's Residue Theorem to do the inversion.