Lecture 3: Jumps

Jim Gatheral, Merrill Lynch*

Case Studies in Financial Modelling Course Notes,
Courant Institute of Mathematical Sciences,
Fall Term, 2002

*I am indebted to Peter Friz for carefully reading these notes, providing corrections and suggesting useful improvements.

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.

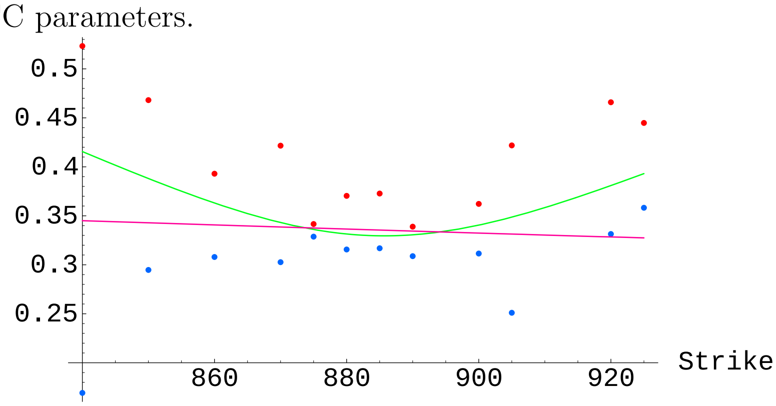

7.2 Derivation of the Valuation Equation

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

$$dS = \mu S\,dt + \sigma S\,dZ + (J-1)S\,dq \tag{40}$$

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

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

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

$$\begin{aligned} \frac{\partial V}{\partial t} &+ \frac{1}{2}\sigma^2\,S^2\frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV \\ &+ \lambda(S,t)\left\{V(JS,t) - V(S,t) - (J-1)S\frac{\partial V}{\partial S}\right\} = 0 \end{aligned} \qquad (41)$$

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 (41) gives

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

Since 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 (41), 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 (41)?

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 can be replicated using stocks and bonds and by extension, option traders can be seen to have social value.

To extend equation (41) to the case of jumps of uncertain size, we need to introduce an additional assumption. The usual assumption due to Merton (1976) is that if jumps are uncorrelated with the market, jump risk is diversifiable and should not be rewarded. Suppose we delta-hedge with Δ of stock (no other assets) with

$$\Delta = \frac{\partial V}{\partial S}$$

Then the change in the hedge portfolio over the time interval dt is given by

$$d\Pi = \left\{ \frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} \right\} dt + \left\{ V(JS,t) - V(S,t) - (J-1) S \frac{\partial V}{\partial S} \right\} dq \tag{42}$$

No reward for jump risk translates into the requirement that the risk-neutral expectation of equation (42) be zero. This gives the following equation for valuing financial assets under jump-diffusion:

$$
\begin{aligned}
\frac{\partial V}{\partial t} &+ \frac{1}{2}\sigma^2 \, S^2 \frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV \\
&+ \lambda(t) \left\{ \mathbf{E}\left[V(JS,t) - V(S,t)\right] - \mathbf{E}\left[J-1\right] S \frac{\partial V}{\partial S} \right\} = 0 \qquad (43)
\end{aligned}
$$

The assumption that jumps are uncorrelated with the market is clearly untenable in the case of the short-dated SPX skew: SPX is the market. Nevertheless, despite this reservation, we will continue with our analysis.

7.3 Lévy Processes

With constant hazard rate λ, the logarithmic version of the jump-diffusion process (40) 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 $\phi_T(u) \equiv \mathbf{E}\left[e^{iux_T}\right]$ has the representation

$$
\phi_T(u) = \exp\left\{ iu\omega T - \frac{1}{2}u^2\,\sigma^2 T + T\int \left[e^{iu\xi} - 1\right]\mu(\xi)\,d\xi \right\} \qquad (44)
$$

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

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

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

Example 1: Black-Scholes

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

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

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

This is the case we are really interested in here. 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 (44), we see that the characteristic function is given by

$$\begin{aligned} \phi_T(u) &= \exp\left\{iu\omega T - \frac{1}{2}u^2\,\sigma^2 T + T\int\left[e^{iu\xi}-1\right]\frac{\lambda}{\sqrt{2\pi\delta^2}}\exp\left\{-\frac{(\xi-\alpha)^2}{2\delta^2}\right\}d\xi\right\} \\ &= \exp\left\{iu\omega T - \frac{1}{2}u^2\,\sigma^2 T + \lambda T\left(e^{iu\alpha - u^2\delta^2/2} - 1\right)\right\} \qquad (45) \end{aligned}$$

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 $\alpha = \delta = 0$.

7.4 Solving the Valuation Equation

Unlike the partial differential equations (PDEs) we are used to solving in derivatives valuation problems, equation (43) 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 methods.

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):

$$C(S,K,T) = S - \sqrt{SK}\frac{1}{\pi}\int_0^\infty \frac{du}{u^2+\frac{1}{4}}\,\mathbf{Re}\left[e^{-iuk}\phi_T\left(u - i/2\right)\right] \qquad (46)$$

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

7.5 Implied Volatility Term Structure and Skew

Equation (46) 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 (46) 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

$$\int_0^\infty \frac{du}{u^2+\frac{1}{4}} \mathbf{Re}\left[e^{-iuk}\left(\phi_T\left(u-i/2\right)-e^{-\frac{1}{2}\left(u^2+\frac{1}{4}\right)\sigma_{BS}^2 T}\right)\right]=0 \tag{47}$$

Equation (47) 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).

7.5.1 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 (47) 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

$$\left.\frac{\partial \sigma_{BS}}{\partial k}\right|_{k=0}=-e^{\frac{\sigma_{BS}^2 T}{8}}\sqrt{\frac{2}{\pi}}\frac{1}{\sqrt{T}}\int_0^\infty du\,\frac{u\,\mathbf{Im}\left[\phi_T(u-i/2)\right]}{u^2+\frac{1}{4}} \tag{48}$$

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

Then

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

To get the very short time behavior of this expression, we can consider the limit $\lambda \to 0$. In that limit, noting that $\mathbf{Im}\left[e^{\psi_T(u-i/2)}\right] = 0$ when $\lambda = 0$, we have

$$\begin{aligned}
\mathbf{Im}\left[e^{\psi_T(u-i/2)}\right] &\approx \lambda \frac{\partial}{\partial\lambda}\mathbf{Im}\left[e^{\psi(u-i/2)T}\right]\Bigg|_{\lambda=0} \\
&= -\lambda T\left\{u\left(e^{\alpha+\delta^2/2}-1\right) + \text{ highly oscillatory terms }\right\}
\end{aligned}$$

Substituting this expression into equation (48), we get

$$\begin{aligned}
\left.\frac{\partial\sigma_{BS}}{\partial k}\right|_{k=0} &= -\mu_J T e^{\frac{\sigma_{BS}^2 T}{8}}\sqrt{\frac{2}{\pi}}\frac{1}{\sqrt{T}}\int_0^\infty du\,\frac{u^2}{u^2+\frac{1}{4}}\,e^{-\frac{1}{2}(u^2+1/4)\sigma^2 T} \\
&\qquad + \text{higher order terms} \\
&\to -\frac{\mu_J}{\sigma} \text{ as } T \to 0
\end{aligned}$$

where $\mu_J = -\lambda(e^{\alpha+\delta^2/2}-1)$ is the adjustment to the risk-neutral drift for jumps.

Intuition

Whilst the above derivation has the virtue of resulting from a systematic approach, a more direct approach proves more intuitive.

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

$$\begin{aligned}
C_J(S,K,\Delta T) &\approx (1-\lambda\Delta T)\,C_{BS}(Se^{\mu_J\Delta T},K,\Delta T) + \lambda\Delta T\,C(JS,K,\Delta T) \\
&= C_{BS}(Se^{\mu_J\Delta T},K,\Delta T) + O(\Delta T) \qquad (49)
\end{aligned}$$

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 (49) 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

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

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

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

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 (49)

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

Then, for small ΔT,

$$\left.\frac{\partial \sigma_{BS}}{\partial k}\right|_{k=0} \approx -\frac{\mu_J}{\sigma_{BS}} \tag{50}$$

as required.

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.

Example 3: Stochastic Volatility plus Stock Price Jumps (SVJ)

Since jumps generate a steep short-dated skew which 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, 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 $\psi(u) = -\lambda_J i u \left(e^{\alpha+\delta^2/2} - 1\right) + \lambda_J \left(e^{iu\alpha - u^2\delta^2/2} - 1\right)$ and $C(u,T)$, $D(u,T)$ are as before.

Again, we may substitute this functional form into equations (47) and (48) to get the implied volatilities and at-the-money volatility skew respectively for any given expiration.

Figure 2 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 3). 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 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 3 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.

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). Not only that, but as we shall see, the extreme short-end skew still can't be explained with reasonable parameters.

7.5.2 Fitting the 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

Figure 2: 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.

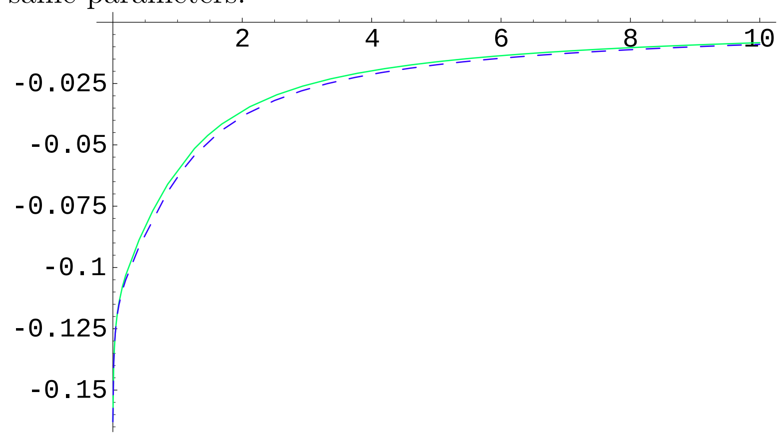

Figure 3: The green line is a graph of the at-the-money variance skew in the SVJ model with BCC parameters vs time to expiration. The dashed red line represents the at-the-money Heston skew with the same parameters.

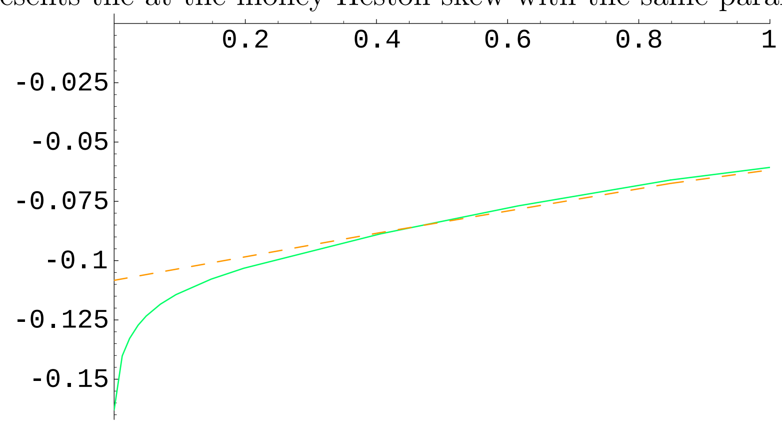

Table 3: Various fits of jump-diffusion style models to SPX data. JD means Jump Diffusion and SVJ means Stochastic Volatility plus Jumps.

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

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 3.

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 SPX volatility skew doesn't really change much over time so it does make some sense to compare them.

We can't help but notice the extreme size of the mean downward jump in the Andersen-Andreasen estimate! It is given by $e^{\alpha+\delta^2/2} - 1 = -54.54\%$. Do we really believe that the market is pricing options as if there is an 8.5% probability per year of a -55% jump in asset prices? Or is it more reasonable to expect a -10% jump with a 40% probability? The implications for option pricing are clearly very different. For example, consider the value of a short-dated option that pays \$1 if the market hits a barrier level set at 50% of the current stock price. For pretty much any sensible volatility, if we believe in average -10% jumps, this option is almost worthless. However, if we believe in -55% jumps, the option has real value given by the probability of a jump over the life of the option.

These crazy parameter estimates simply reflect the fact that the jump-

Figure 4: Graph of SPX ATM variance skew vs. time to expiration in the jump-diffusion model with Andersen-Andreasen (orange line) and the SVJ model with Matytsin parameters (blue line). The green points represent empirical variance skews from the Andersen-Andreasen paper and the red points empirical variance skews from September 17, 2002.

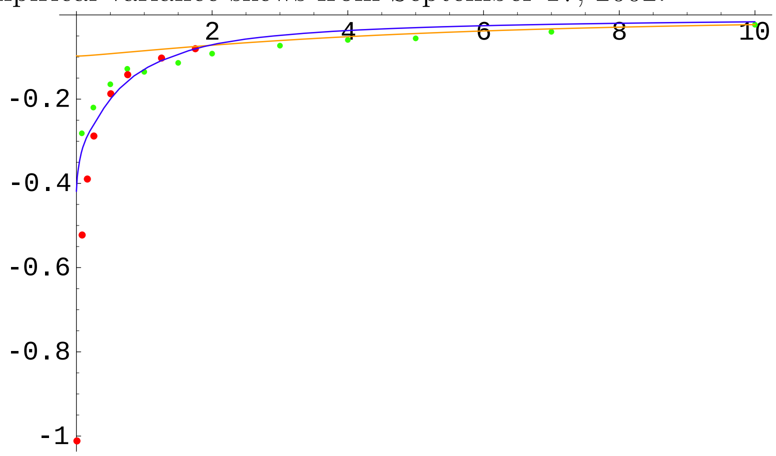

diffusion model is totally misspecified: a closer analysis reveals that volatility surface implied by a jump-diffusion model just doesn't look like the empirically observed implied volatility surface.

7.5.3 Fits versus Reality

Figure 4 plots the at-the-money variance skews corresponding to the Andersen-Andreasen and Matytsin fits and superimposes empirically observed skews from September 17, 2002 and from the Andersen-Andreasen paper. Just looking at the graph, we see that the SVJ model with Matytsin parameters fits better than the jump-diffusion model with Andersen-Andreasen parameters. Also, we note that the short-dated skews from September 17, 2002 are pretty consistent with longer-dated skews from the Andersen-Andreasen paper and so also consistent with our claim that volatility skew is relatively stable over time.

Most strikingly however, just looking at the empirical skews, we see that

$$\left.\frac{\partial \sigma_{BS}^2}{\partial k}\right|_{k=0} \to -1.1 \text{ or so as } t \to 0$$

and neither the JD nor the SVJ models fit even approximately in that limit.

From equation (50), in the JD model, the *variance* skew at $t = 0$ should be given by

$$\left.\frac{\partial \sigma_{BS}^2}{\partial k}\right|_{k=0} \approx -2\mu_J = 2\lambda\left(e^{\alpha+\delta^2/2} - 1\right)$$

and in the SVJ model, there would be an additional term $\rho\eta/2$. In fact, it's easy to see that there is no sensible choice of jump-diffusion parameters that could fit the short-dated skew. To achieve that we would need

$$2\lambda\left(e^{\alpha+\delta^2/2} - 1\right) = 2\lambda\left(\mathbf{E}[J] - 1\right) \approx -1$$

For example, it is easy to see that we must have $\lambda \geq 0.5$. If we do put $\lambda = 0.5$, we get $\mathbf{E}[J] = 0$. This corresponds to a 40% chance per annum of the stock market going to zero! Alternatively, we can set $\mathbf{E}[J]$ to a more reasonable 0.9 to find $\lambda = 10$ which corresponds roughly to a -10% jump every month.

7.5.4 Skew Decay in a Jump-Diffusion Model

The volatility skew decays very rapidly in a jump-diffusion model beyond a certain time to expiration. To estimate this characteristic time, 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 by equating

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

With Bakshi, Cao and Chen parameters, we obtain that the characteristic time beyond which we expect the skew to be flat is around 2 months (consistent with the extra blip in the skew in Figure 3) and with Andersen-Andreasen parameters, the characteristic time is over 9 years, consistent with Figure 4. Now we see why Andersen and Andreasen have that crazy jump size estimate: without that their skew would decay far too fast to remotely fit the implied volatility surface even for longer-dated options.

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

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". Interestingly, as noted by Matytsin (1999) for example, this extension also generates extreme short-dated implied volatility skews consistent with observation.

7.7 Testing for Jumps

A recent paper by Carr and Wu (2002) studies the behavior of SPX option prices as time to expiration approaches zero and concludes that there is strong evidence in the options market for the existence of a jump component in the risk-neutral price process. They conclude that there must also be jumps in the statistical (historical) price process (otherwise Girsanov's Theorem would induce no jumps in the risk-neutral process).

8 Some Applications of Jumps

8.1 Merton's Model of Default

As we have come to expect, Wilmott (1998) gives an excellent introduction to the modelling of default risk. There are two broad types of default-risk model used by practitioners: so-called *structural* models and so-called *reduced form* models. I found the following useful description by JabairuStork on Wilmott.com:

"A *structural model* (of firm default) postulates that default occurs when some economic variable (like firm value) crosses some barrier (like debt value), typically using a contingent claims model to support this assertion and to find the probability of default. Both H-W and Creditgrades are models of this form."

"A *reduced form* model models default as a random occurrence - there is no observable or latent variable which triggers the default event, it just happens. The Duffie-Singleton model is a reduced form model. These models are easy to calibrate, but because they lack any ability to explain why default happens, I think they make most people nervous. Basically, you estimate an intensity for the arrival of default (possibly as a function of time, possibly as a stochastic process, possibly as a function of other things.)"

Merton's model is the simplest possible example of a reduced form model. It supposes that there some probability $\lambda(t)$ per unit time of the stock price jumping to zero (the *hazard rate*) whereupon default occurs. Jumps are independent of the stock price process. Then, contingent claims must satisfy equation (43) with $\mathbf{E}[J] = 0$. It is particulary straightforward to value a call option because for a call, $V(SJ, t) = 0$. Substitution into equation (43) gives:

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV - \lambda(t)\left\{V - S\,\frac{\partial V}{\partial S}\right\} = 0 \qquad (51)$$

We immediately recognize equation (51) as the Black-Scholes equation with a shifted interest rate $r + \lambda$. Its solution is of course the Black-Scholes formula with this shifted rate.

The meaning of this shifted rate is particularly clear if we assume no recovery (in the case of default) on the issuer's bonds so that $B(JS, t) = 0$. Then, the risky bond price $B(t, T)$ must also satisfy equation (51) with the solution

$$B(t, T) = e^{-\int\limits_t^T (r(s)+\lambda(s))ds}$$

We identify the shifted rate $r + \lambda$ with the yield (risk-free rate plus credit spread) of a risky bond. The situation is a little more complicated (but not too much more) if we allow some recovery R on default.

Intuition

It may at first seem surprising that the Black-Scholes formula could be a solution of an equation that has a jump to zero (the so-called *jump to ruin*) in it. There is an economic reason for this however.

Recall that the derivation of the Black-Scholes formula involves the construction of a replicating portfolio for a call option involving just stock and risk-free bonds. Suppose instead, we were to construct this portfolio using stock and risky bonds. So long as there is no jump to ruin, the derivation goes through as before and the portfolio is self-financing. If there is a jump to ruin, assuming no recovery on the bond, both the bond and the stock jump to zero – the portfolio is still self-financing!

Figure 5: 3 month implied volatilities from the Merton model assuming a stock volatility of 20% and credit spreads of 100bp (orange), 200bp (green) and 300bp (magenta).

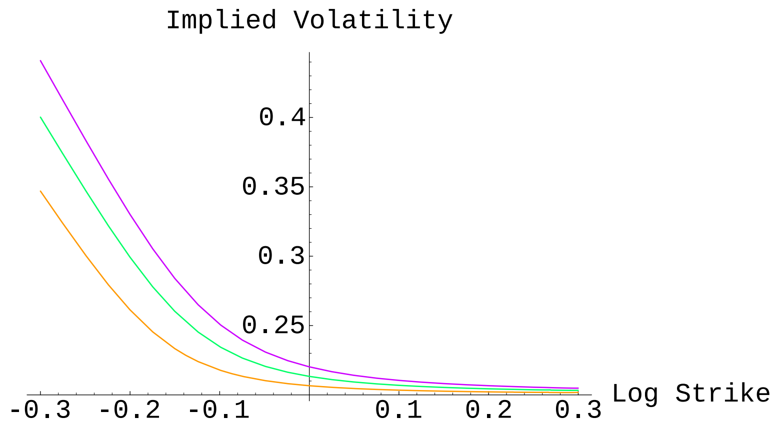

What would happen if we were to hedge a short call option position using stock and *risk-free* bonds following the standard Black-Scholes hedging recipe (as most practitioners actually do)? We would be long stock and short risk-free bonds and in the case of default, the call would end up worthless, the stock would be worthless and we would get full recovery on our risk-free bonds. In other words, on default, we would have a windfall gain. On the other hand, relative to hedging with risky bonds, we would forego the higher carry (or yield).

Implications for the Volatility Skew

All issuers of stock have some probability of defaulting. There is a very active credit derivative market (see DefaultRisk.com for background) which prices default-risk. Black-Scholes implied volatilities are computed by inserting the risk-free rate into the Black-Scholes formula. However, as we just showed, in Merton's model, call option prices are correctly obtained by substituting the risky rate into the Black-Scholes formula. This induces a skew which can become extremely steep for short-dated options on stocks whose issuers have high credit spreads.

In Figure 5, we graph the implied volatility for various issuer credit spreads assuming that options are correctly priced using the Merton model. We see that the downside skew that the model generates can be extreme.

8.2 Capital Structure Arbitrage

Capital structure arbitrage is the term used to describe the current fashion for arbitraging equity claims against fixed income and convertible claims. At its most sophisticated, practitioners build elaborate models of the capital structure of a company to determine the relative values of the various claims - in particular, stock, bonds and convertible bonds. At its simplest, the trader looks to see if equity puts are cheaper than credit derivatives and if so buys the one and sells the other. To understand this, we review put-call parity.

Put-Call Parity

We saw above that in the Merton model, the value of an equity call option is given by the Black-Scholes formula for a call with the risk-free rate replaced by the risky-rate. What about put options? To make the arguments above work, the put option would need to be worthless after the jump to ruin occurs. That would be the case if the put in question were to be written by the issuer of the stock. In that case, when default occurs, assuming zero recovery, the put options would also be worth nothing. So the Black-Scholes formula for a put with the risk-free rate replaced by the risky-rate does value put options written by the issuer.

What about put options written by some default-free counterparty (for example an exchange)? When default occurs, this put option should be worth the strike price. We already know how to value a call written by a default-free counterparty; by definition, the issuer of a stock cannot default on a call on his own stock so the value of a call written by the issuer of the stock equals the value of a call written by a default-free counterparty. We obtain the value of a put by put-call parity: using risk-free bonds in the case of the default-free counterparty and risky bonds in the case of the risky counterparty.

Denoting the value of a risk-free put, call and bond by P_0,C_0 and B_0 and the value of risky claims on the issuer of the stock by P_I, C_I and B_I (I for issuer), we obtain

$$\begin{aligned}
P_0 &= C_0 + KB_0 - S \text{ (from put-call parity with risk-free bonds)} \\
&= C_I + KB_0 - S \text{ (risk-free and issuer-written calls have the same value)} \\
&= P_I + S - KB_I + KB_0 - S \text{ (from put-call parity with risky bonds)} \\
&= P_I + K(B_0 - B_I)
\end{aligned}$$

As we would expect, the risk-free put is worth more than the risky put. The excess value is equal to the difference in risky and risk-free bond prices (times the strike price). With maturity-independent rates and credit spreads for clarity and setting $t = 0$, we obtain

$$B_0 - B_I = e^{-rT}\left(1 - e^{-\lambda T}\right)$$

which is just the discounted probability of default in the Merton model. In words, the extra value is the strike price times the (pseudo-) probability that default occurs. This payoff is also more or less exactly the payoff of a default put in the credit derivatives market.

The Arbitrage

Referring back to Figure 5, we see that the downside implied volatility skew can be extreme for stocks whose issuers have high credit spreads. Equity option market makers (until recently at least) made do with heuristic rules to determine whether a skew looked reasonable or not; implied volatility skews of the magnitude shown in Figure 5 seemed just too extreme to be considered reasonable. Taking advantage of the market maker's lack of understanding, the trader buys an equity option on the exchange at a "very high" (but of course insufficiently high) implied volatility and sells a default put on the same stock in the credit derivatives market locking in a risk-free return.

8.3 The Baseball Trade

The baseball trade takes its name from the familiar "three strikes and you're out" rule. It is one of the most model-dependent structures ever traded. Its terms are as follows:

- We establish an initial range for the stock price (95 to 105 say).
- If the stock price has exited the range at any reset point, we reset the barriers to be equidistant from the *new* stock price.
- We repeat this procedure until the third time the stock price exits a range whereupon the trade expires worthless.
- If the stock price is still within one of the specified ranges at maturity, the trade pays \$1.

In fact, this structure was popular for a while in the FX markets where jumps are uncommon. In equity markets however, where jumps are common, the extreme model-dependence is poisonous.

To see this, suppose we were to value it using a jump-diffusion model with Andersen-Andreasen parameters as in Table 3. The diffusion component of the process would have rather low volatility and the probability of having more than one jump in one year say is negligible. If and when a jump occurs, the new range is set around the point reached *not* around the old range. If the original range is sufficiently wide the probability of this claim paying \$1 at the end is very high.

On the other hand, if we value the same trade using a local volatility model where the local volatilities are calibrated to return jump-diffusion model European option prices, the volatility will be that much higher and the probability of not getting \$1 that much greater.

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 Madan, 1999, Option valuation using the fast fourier transform, *Journal of Computational Finance* 2, 61–73.

Carr, Peter, and Liuren Wu, 2002, What type of process underlies options? a simple robust test, Discussion paper Courant Institute, New York University.

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.

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. 26, pp. 325–335 (John Wiley & Sons: Chichester).

A Proof of Equation (46)

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

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

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.