16

Fear of Jumps[1]

16.1 Introduction

Will there ever be a generally accepted option pricing model? One that academics will teach their students, and practitioners will use to trade? I don't know the answer, but I have opinions about things we'll find in model X.

It's obvious that model X must account for any economically important behavior that we find in security prices. Stock price changes are largely unpredictable in sign (little raw auto-correlation) at most time scales, so we need a model with that. Stock price volatility, which can be defined in many ways, is persistent but not constant: it has a very significant random variation. Briefly, model X needs stochastic volatility.

For example, at this writing, we are a few weeks past a month of extreme volatility in the equity markets (July, 2002). The major (U.S.) market indices were roughly twice as volatile as normal. There were many 3 to 4% daily moves in the broad market. The CBOE's VIX index (a measure of expected volatility of the S&P 100 Index[2]) fluctuated between 40% and 50% over most of the second half of the month. The reasons? Well, one can always point to events a build-up of fear and anxiety after a 2 1/2 year bear market, driven by over-valuation and collapse, with a threat of renewed terrorism and a collection of infamous corporate frauds thrown in the mix – but, in reality, no one really knows why these events trigger this reaction at this particular time. Of course, next time it will all be different, but surely stochastic volatility will always be with us.

For most of this recent bear market, it's been a pretty orderly sell-off. The academics would call it a continuous sample path. The exception, of course, was the September 11 World Trade Center attack. As witnesses to a mass murder, we were profoundly shocked. However, we don't need to invoke extreme horror as a prototype for market jump triggers. Indeed, internal market events of transitory historical importance and devoid of criminality (such as the Oct '87 market crash) provide better examples. Some mini-crashes could be invoked. And, below these systematic events, down at the micro level, consider the panoply of run-of-the-mill corporate events that regularly trigger trading halts and subsequent jumps. Our conclusion: discontinuous price moves are another ingredient for which model X must account. And that is our subject this month: option valuation methods that handle jumps. I'll write exclusively about Euro-style options – we'll consider American-style options next time.

[1] First published in Wilmott Mag., (Dec., 2002), 60-67. This version: minor revisions.

[2] The CBOE persists in applying a strange time scaling in the VIX calculations. This anomaly was first noted by Stephen Taylor of Lancaster Univ. The result is that if you work out your own at-the-money implied volatility for the OEX, you will almost certainly get a smaller number than the VIX. Be warned! (2015 note added: the remark applies to the 'old VIX'; the methodology was substantially revised in 2003).

Fig. 16.1. SPX Options: Implied Volatility vs. Strike

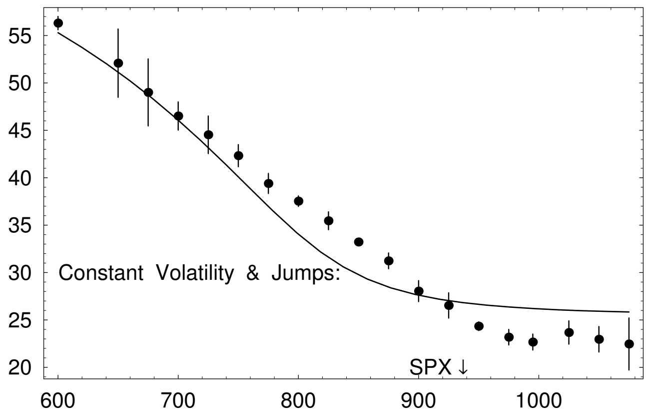

Notes. Data from August 16, 2002 (1 month to expiration)

Fig. 16.2. SPX Options: Implied Volatility vs. Strike

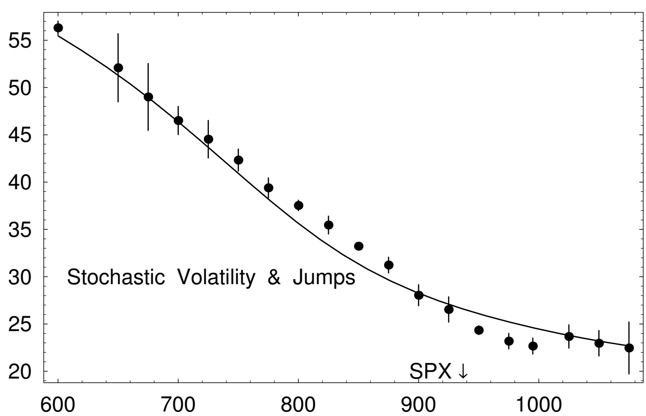

Notes. Data from August 16, 2002 (1 month to expiration)

16.2 Market snapshot

Fear of a market-wide jump is a natural explanation for pricing effects in short-dated, broad-based, index options, like the S&P500. For example, at this writing, the next SPX option expiration is about a month away. By converting market option prices into Black-Scholes prices, the result is an 'Implied Volatility Smile' (Figure 16.1). We used closing option quotes for the September '02 expiration. Error bars are drawn from the implied volatility at the option bid to the ask.

The solid line drawn in the figure, labeled 'Constant Volatility + Jumps', is a rough fit to the market smile, using Merton's jump-diffusion model. This model, discussed below, adds a Poisson-driven jump process (with log-normally distributed jumps) to the Black-Scholes model. There are three extra parameters. First, there is the mean jump event arrival frequency λ. I set $\lambda = 0.30$/year which says that the jumps (in a 'risk-adjusted world') arrive once every 3 1/3 years on average. Next, when the index price S jumps by ΔS , there is the mean value for $\log[(S + \Delta S)/S]$, called μ_J . I set $\mu_J = -0.25$, which says that the mean logarithmic jump amplitude is -25%. In fact, when a jump occurs, the actual value of $\log[(S + \Delta S)/S]$ is random and is drawn from a normal distribution. This normal distribution has mean μ_J and volatility σ_J. I set $\sigma_J = 0.15$, which means that, when there is a jump, the logarithmic jump amplitude volatility is 15% about the mean jump size.[3]

This smile chart has a typical SPX shape. The overall volatility levels are above-average because we are coming off the extreme July '02 volatility mentioned earlier. The nice thing about jump models is their easy fit to very steep smile charts. It's easy because, with Poisson-driven jumps, the smile actually becomes infinitely steep as expiration approaches.

Now you may object and say that for the market to expect a -25% crash every 3 years or so is not realistic: it's too paranoid. Our answer is: you're right, but these parameters don't say that. These parameters actually represent a mathematical combination of (i) what the market expects in the way of possible crashes, plus (ii) the market's subjective fear of the event, as reflected by its risk aversion. The parameter values that I quoted can be compatible with a crash event actually expected to occur, say once every 10 years, on average, and of lesser magnitude when it does occur. Confused? Then read on, because by Section 16.4, we explain this apparent paradox.

Before getting into the mathematics of jumps, an important shortcoming of jump models is visible in the smile chart and worth noting. An arrow in the chart that shows the current value of the SPX index. Look what happens to the smile for the strikes above that value. As the strikes increase to the right, the market implied volatilities keep declining, although not as rapidly as on the left. But the theoretical model (the solid line) misses this behavior because it is flattening out rapidly to the right. The reason the theoretical model is flattening is that our jump distribution has little effect on these (upper strikes), since the jumps are concentrated about -25% (although with some spread). So, where the jumps have negligible effect, we are back to the Black-Scholes model, which has a flat smile. But, the market has other ideas. The market is 'correct' because what typically happens is that, if the index price moves up from its current level, then volatility will decline. This is the well-known 'leverage' effect. (It also happens in reverse for price declines and helps exaggerate the smile steepness on the left). This leverage behavior is easily captured by stochastic volatility models, but is missing from the jump-diffusion models. Figure 16.2 shows the same market smile, but where we allow stochastic (diffusion) volatility in addition to the jumps for the theoretical comparison model.

For Figure 16.2, I added a GARCH diffusion volatility process to the mix. I'm not going to go into details because it's tangential to this month's topic. The only new parameter worth noting is my stock-volatility correlation: $\rho = -0.65$, which is typical of SPX estimates. As expected, adding stochastic volatility improves the theoretical fit because now the implied volatilities keep declining as the strikes go up (at least on the range of the chart). So Figure 16.2 emphasizes the point made earlier: we need both stochastic volatility and jumps in the ultimate model X. Now let's go on and discuss some jump model mathematics.

[3] In addition, there are the Black-Scholes parameters: I used an interest rate of 1.8%, a dividend yield of 1.7%, and a (diffusion) volatility of 25%. I didn't do any fancy optimization to get these parameters, but just tried a few fairly round numbers that I knew from experience would be decent fits.

16.3 Merton's jump-diffusion model

First, a little setup. In the Black-Scholes theory, stock prices follow geometric Brownian motion (GBM). The objective (actual) stock price process is given by $P: \; dS_t/S_t = \alpha dt + \sigma dB_t$, where P is a label for the real-world. Here α is a constant drift rate, σ is the diffusion volatility, and B_t is a standard P-Brownian motion. Equivalently, $S_t = S_0 \exp X_t$, where $X_t = (\alpha - \sigma^2/2)t + \sigma B_t$. In this article, S_t is an index price, representing the entire equity market. Because of that, α is the expected total return for stocks in the aggregate (including dividends for simplicity).

A compelling feature of this model is *level-independence* – we have the same theory at Dow 1000 or Dow 10,000. In our generalizations, I preserve this. Also, since our stock index is a proxy for market wealth, we want level-independence for that also: the same theory if the market capitalization is \$1 trillion dollars or \$10 trillion dollars, etc. In my thinking, if you don't like level-independence, which seems very natural, then it's your burden to explain why particular index or capitalization levels are important (because these particular levels will always play a special role in your alternative theory).

Merton [186] adds stock price jumps in this level-independent way – he does so by adding a Poisson-driven jump process, with multiplicative jump amplitudes Y_t. That is, when there is no jump, the stock evolution still follows GBM. But, when a jump occurs at time t, he postulates $S_t = S_{t-} \times Y_t$, where S_{t-} is the stock price just before the jump and Y_t, although randomly drawn, has no dependence on S_{t-}.

If you are new to jump models, it may be helpful if you think of a two-step process. In step one, we see if a jump event has occurred – this might be thought of as the news event associated with the jump. Second (but really simultaneously), the stock market reacts by deciding how much it will jump (and jumping!) Mathematically, this is described as follows.

Jump events arrive randomly; their arrival is described by a Poisson process with mean arrival rate λ. Under a Poisson process, during a small time interval Δt , the probability of a jump event occurring is $\lambda \Delta t + O(\Delta t^2)$, independent of anything that has happened before. Hence, the probability of no jump is $1 - \lambda \Delta t + O(\Delta t^2)$. Technically, Merton's jump process is called a *compound* Poisson process because, if a jump occurs, then its magnitude Y_t is determined by making a completely independent random drawing from another probability distribution.

It's actually much more convenient to write $Y = Y(x) = \mathrm{e}^x$, where x is the logarithmic jump amplitude, which we draw from some probability distribution $p(x)$. Throughout, x ranges from $-\infty$ to $+\infty$, so when you see an integral over x, those are the ranges. As we stress, $p(x)$ is independent of the stock price. As an explicit SDE, we have

$$\frac{dS_t}{S_t} = \begin{cases} (\alpha - \lambda k)dt + \sigma dB_t, & \text{when no jump occurs,} \\ (\alpha - \lambda k)dt + \sigma dB_t + (\mathrm{e}^x - 1), & \text{when a jump occurs and } x \sim p(x). \end{cases} \tag{16.1}$$

The notation $x \sim p(x)$ means that the random variable x is drawn from the probability distribution $p(x)$. In (16.1), the constant $k = \int (\mathrm{e}^x - 1)p(x)dx$ is introduced to preserve the meaning of α as the expected total rate of return. This real-world theory can be summarized by four quantities: (i) the stock expected rate of return α, (ii) the Brownian volatility scale σ, the mean jump arrival rate λ, and (iv) the jump amplitude distribution $p(x)$. A shorthand for the objective process is P: $\{\alpha, \sigma, \lambda, p(x)\}$.

Readers of this magazine will know that even if we have discovered a perfect description of the stock price evolution in (16.1), we still can't price an option until we know how to 'risk-adjust' that evolution process. We will call the risk-adjusted process Q. In the (non-jumping) Black-Scholes theory, with P: $\{\alpha, \sigma\}$, the transformation is unique: P$\rightarrow$Q: $\{r, \sigma\}$. Again, this is shorthand to indicate that, under Q: $dS_t/S_t = rdt + \sigma dB_t^Q$, where $B^Q(t)$ is a Q-Brownian motion. The stock price drift parameter has changed $\alpha \to r$, where r is a (constant) interest rate.

Unfortunately, once jumps are introduced, many risk adjustments are possible in principle. Which one does the market choose? Merton postulated 'non-systematic' jumps, a theory of firm-specific (idiosyncratic) jumps. It's really the opposite of what we are discussing, since we are considering market-wide, systematic jumps. Nevertheless, it's a good starting point for our discussion. With his assumption and our short-hand, he derives P: $\{\alpha, \sigma, \lambda, p(x)\} \to$ Q: $\{r, \sigma, \lambda, p(x)\}$.

That is, in Merton's theory, the jump process remains exactly the same after risk adjustment. The only change is that the equity expected return becomes r. Written out, his theory is that

$$\text{Under Q: } \frac{dS_t}{S_t} = \begin{cases} (r - \lambda k)dt + \sigma dB_t^Q, & \text{when no jump occurs,} \\ (r - \lambda k)dt + \sigma dB_t^Q + (\mathrm{e}^x - 1), & \text{when a jump occurs and } x \sim p(x). \end{cases} \tag{16.2}$$

Then, European-style option prices, say a call option $C(S, t)$, must satisfy an equation obtained by requiring that $\exp(-rt)C(S, t)$ is a Q-martingale. This equation generalizes the Black-Scholes partial differential equation (PDE). Technically, the new equation is an 'integro-PDE'. Using subscripts to indicate differentiation, he found:

$$C_t - rC + (r - \lambda k)SC_S + \frac{1}{2}\sigma^2 S^2 C_{SS} + \lambda \int [C(S\mathrm{e}^x, t) - C(S, t)]\, p(x)dx = 0. \tag{16.3}$$

The terminal condition is: at expiration $t = T$, $C(S, T) = \max(S - K, 0) = (S - K)^+$, where K is the strike.

Merton discovered he could solve (16.3) when $p(x)$ is a normal distribution; i.e., $p(x) = \exp\{-(x - \mu_J)^2/2\sigma_J^2)\}/(2\pi\sigma_J^2)^{1/2}$. With that specialization, μ_J is the mean logarithmic jump amplitude and σ_J is the jump volatility. This specialized case is often what people mean when they say *the* jump-diffusion model. His solution relied on the special property that a product of log-normal variates also has a log-normal distribution. His solution was an infinite sum of Black-Scholes functions:

$$C(S, K, t, T) = \sum_{n=0}^{\infty} \exp(-\hat{\lambda}\tau) \frac{(\hat{\lambda}\tau)^n}{n!}\, c(S, K, \tau, \nu_n^2, r_n), \tag{16.4}$$

where $\tau = T - t$, $\hat{\lambda} = \lambda(1 + k)$, $\nu_n^2 = \sigma^2 + n\sigma_J^2/\tau$, and $r_n = r - \lambda k + n(\mu_J^2 + \sigma_J^2/2)$. In (16.4), the 'small c' functions are the Black-Scholes solution functions

$$c(S, K, \tau, \sigma^2, r) = S\,\Phi(d_+) - \mathrm{e}^{-r\tau} K \Phi(d_-),$$

$$d_\pm = \frac{1}{\sigma\sqrt{\tau}} \left[\log\left(\frac{S}{\mathrm{e}^{-r\tau} K} \right) \pm \frac{1}{2}\sigma^2 \tau \right],$$

and

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} \exp(-y^2/2)\, dy.$$

16.4 Better risk adjustment

For market-wide jump effects we need a better risk adjustment. My favorite approach is a utility-based equilibrium theory. This is an approach taken by Bates [20, 21], Naik and Lee [195], and others. In the simplest version, you postulate a representative investor who invests to maximize her utility of (terminal) wealth $U(W) = W^\gamma$, where γ is a constant risk-aversion parameter. When $\gamma = 1$, we say that investors are risk-neutral because they will price stocks so that $\alpha = r$ under the *objective* process, not just the risk-adjusted one. When $\gamma < 1$, we say that investors are strictly *risk-averse*, which is what we really expect. Our results below will hold for any $\gamma \le 1$. Besides level-independence, power utility is known to be particularly compatible with the Black-Scholes' theory — see Bick [24]. In addition, lots of man-hours have been invested by economists trying to estimate γ. While just about any risk-adjustment approach has an ad-hoc air about it, I find this one compelling because it brings in risk aversion as a single new constant parameter. One small step at a time ...

Since our stock price S is an index for market wealth W, then these two random variables are proportional and so have the same evolution equations. To apply the theory, you allow the representative to form a portfolio potentially consisting of a stock index fund, a certain amount of borrowing or lending at the rate r, and an option on the index fund. But, the representative's portfolio is the aggregate of all the individuals in the market. Both borrowing and lending, and long and short option positions, are assumed to be contracts between individuals. In the aggregate, there can be no net borrowing and no net option position – 100% of the market wealth is invested in the index fund. So, the maximizing strategy for the representative must also be one that clears the market in this sense. These two "market clearing" conditions yield two relations: (i) an equation that determines α, and (ii) an equation that determines the option price. I'm not going to show the details, but just report the results. The equation for the Euro-style option price turns out to be

$$C_t - rC + (r - \lambda b)SC_S + \frac{1}{2}\sigma^2 S^2 C_{SS} + \lambda \int [C(S\mathrm{e}^x, t) - C(S,t)]\, p(x)\, \mathrm{e}^{(\gamma-1)x} dx = 0, \tag{16.5}$$

where $b = \int (\mathrm{e}^x - 1)\mathrm{e}^{(\gamma-1)x} p(x) dx$. Notice how similar this is to (16.3). By comparing the integral terms, you would naturally want to introduce a new jump amplitude distribution $q(x) = p(x)\,\mathrm{e}^{(\gamma-1)x} / \int p(x)\mathrm{e}^{(\gamma-1)x} dx$. This can be achieved by introducing a rescaled Poisson arrival rate $\lambda^Q = \lambda \int p(x)\mathrm{e}^{(\gamma-1)x} dx$. Finally, we observe that $\lambda b = \lambda^Q k^Q$, where $k^Q = \int (\mathrm{e}^x - 1) q(x) dx$. So, our option pricing equation becomes

$$C_t - rC + (r - \lambda^Q k^Q)SC_S + \frac{1}{2}\sigma^2 S^2 C_{SS} + \lambda^Q \int [C(S\mathrm{e}^x, t) - C(S,t)]\, q(x)\, dx = 0. \tag{16.6}$$

In other words, under our equilibrium theory for the risk adjustment, we still have a combination of a Brownian motion with drift and a Poisson-driven jump process. However, the drift parameter has changed, the Poisson arrival rate has changed, and the jump amplitude distribution has changed. To summarize, in our short-hand, we find the following transformation rules for broad-based index options:

Under power utility equilibrium: $U(W) = W^\gamma$;

then, $\quad P: \{\alpha, \sigma, \lambda, p(x)\} \to Q: \{r, \sigma, \lambda^Q, q(x)\}$,

$$\text{where} \quad \lambda^Q = \lambda \int p(x)\mathrm{e}^{(\gamma-1)x} dx, \quad \text{and} \quad q(x) = \frac{p(x)\mathrm{e}^{(\gamma-1)x}}{\int p(x)\mathrm{e}^{(\gamma-1)x} dx}. \tag{16.7}$$

Even more explicitly, the risk-adjusted stock price evolution is given by

$$\text{Under Q:} \quad \frac{dS_t}{S_t} = \begin{cases} (r - \lambda^Q k^Q)\, dt + \sigma dB_t^Q, & \text{when no jump occurs,} \\ (r - \lambda^Q k^Q)\, dt + \sigma dB_t^Q + (\mathrm{e}^x - 1), & \text{when a jump occurs; } x \sim q(x). \end{cases} \tag{16.8}$$

If you accept (16.8), then (16.6) is easy to derive without utility theory. Again, you just require that $e^{-rt}C(S,t)$ be a Q-martingale. Applying Itô's lemma (generalized for jumps), yields (16.6). In fact, there is a general theory (see Colwell and Elliott [57]) of possible equivalent martingale measure transformations for Poisson-driven jumps; this general theory also leads to (16.7) except that the specific utility term $e^{(\gamma-1)x}$ becomes an arbitrary multiplicative $g(x)$. What the utility maximization adds to the general theory is a specific transformation rule connecting $p(x)$ and $q(x)$.

To get our option price formula, we need to 'solve' (16.8) for the (risk-adjusted) stock price evolution. Using the new drift parameter

$$\omega = r - \frac{1}{2}\sigma^2 - \lambda^Q k^Q = r - \frac{1}{2}\sigma^2 - \lambda^Q \int (e^x - 1)\, q(x)\, dx, \tag{16.9}$$

the solution to (16.8) is

$$S_T = S_0 \, \exp\left\{\omega T + \sigma B(T)\right\} \prod_{0<t\leq T} e^{x_t} 1_{\{N_t \neq N_{t-}\}}. \tag{16.10}$$

In (16.10), $N_t = 0, 1, 2\ldots$ is the cumulative number of jumps that have occurred between time 0 and t. The capital 'Pi' is a product, and is just a fancy way of saying: there is a jump at this time t because the cumulative number of jumps has changed. So, all this term is doing is, at each jump time, multiplying the Brownian motion part of the solution by the jump amplitude e^x. If there is no jump at time t, then the product term is defined to be 1. (not zero!) While this notation may seem awkward, it's really just a restatement of the Poisson jump model, except now based upon λ^Q and $q(x)$. The beauty of this notation is that it's easy to take S_T to a power, say $(S_T)^p$. Then, we can write immediately

$$(S_T)^p = (S_0)^p \, \exp\left\{p\,\omega T + p\,\sigma B(T)\right\} \prod_{0<t\leq T} e^{p x_t} 1_{\{N_t \neq N_{t-}\}}, \tag{16.11}$$

a relation that will be important further on.

16.5 Risk adjustment examples

Example I : Log-normal jumps revisited.

Now we look at Merton's model in a new light. Under our risk-adjustment,

$$q(x) \propto e^{(\gamma-1)x} p(x) \propto e^{(\gamma-1)x} \exp\left[-\frac{(x-\mu_J)^2}{2\sigma_J^2}\right] \propto \exp\left[-\frac{(x-\mu_J^Q)^2}{2\sigma_J^2}\right],$$

where we are freely dropping proportionality constants, and $\mu_J^Q = \mu_J - (1-\gamma)\sigma_J^2$. This example illustrates that the log-normal jump case is special, in the sense that even after the power utility risk-adjustment, we still have a log-normal jump amplitude distribution. In general, if $p(x)$ is from a well-known 'class' then after you multiply by the risk adjustment, there is no particular reason to still be in the same class. (See the next example). The risk-adjusted Poisson arrival rate is given by

$$\lambda^Q = \lambda \int p(x)\, e^{(\gamma-1)x}\, dx = \lambda \, \exp\left\{(\gamma-1)\mu_J + \frac{1}{2}(\gamma-1)^2\sigma_J^2\right\}.$$

For a numerical example, suppose your estimate is that, statistically, market jumps occur once every 10 years on average ($\lambda = 0.10$), with a mean logarithmic jump of 25%

($\mu_J = -0.25$) and volatility of 15% about that ($\sigma_J = 0.15$). Take a modestly negative value for the risk-aversion parameter, say $\gamma = -1.5$. Then, after risk adjustment, we have $\lambda^Q \approx 0.20$, $\mu_J^Q \approx -0.31$. So options on the index will be priced as if the jumps occur once every 5 years on average, and are more negative on average than the actual process. The stronger the risk aversion, the stronger the effect.

This numerical example illustrates a number of important points about stock price jumps. First, it shows that jump distributions can be *very sensitive* to modest risk adjustments. The possibility of a large negative market jump becomes exaggerated by the risk aversion in two ways: toward a higher frequency of occurrence and toward a more negative outcome. So the risk adjustment can resolve the apparent paradox of Section 16.2, where the market implied jump parameters seemed too extreme.

Another point. Consider that the total (experienced) volatility of the stock price process, (under either measure) is a sum of two prices: the Brownian diffusion volatility and the jump process. That is, $\sigma_{total}^2 = \sigma^2 + \lambda(\mu_J^2 + \sigma_J^2)$. Now, a lesson from the (non-jumping) Black-Scholes model is that volatility is not changed by risk adjustment. The jump model shows that this lesson is generally wrong. We see that the risk-adjusted total volatility in the jump-diffusion model, $\sigma_{total,Q}^2 = \sigma^2 + \lambda^Q(\mu_{J,Q}^2 + \sigma_J^2)$, can be significantly higher than the actual volatility. So the volatility that options price is subjective. We could have titled this article "Why the market is more volatile than you guessed".

Finally, we see significant implications for any statistical analysis of market-wide jumps. The resulting parameter estimates cannot be used directly for broad-based option pricing. They need risk corrections.

Example II: Kou's double exponential jump model.

With this model, Steve Kou [152] postulates a Poisson-driven jump process, but he starts directly with the risk-adjusted Q-process. Specifically, he takes λ^Q as given and $q(x) = \exp(-|x-\kappa|/\eta)/(2\eta)$, where κ and η are two constants, with the restriction that $0 < \eta < 1$. By relying on some special properties of the exponential distribution, he develops the option solution to (16.6), a solution which is quite complicated although not too hard to code. Below, we show an easier alternative formula.

For right now, we only note that our procedure allows one to work backward to the historical (objective) jump distribution. This model plus our risk-adjustment implies that $p(x) \propto \exp[(1-\gamma)x - |x-\kappa|/\eta]$. Note that $p(x)$ *cannot* be written in the form $p(x) = \exp(-|x-\hat{\kappa}|/\hat{\eta})/(2\hat{\eta})$, for some constant parameters $\hat{\kappa}$ and $\hat{\eta}$. That is, $p(x)$ is *not* of the same double exponential type – not necessarily a bad thing. The model is still easy to work with (see below) and may indeed work better than the Merton model. Moreover, since rare jumps are so ... well, rare ... maybe it makes more sense to start with the risk-adjusted expectations. The main lesson of this example is that objective and risk-adjusted distributions will typically *not* be members of the same parameterized class.

Example III: A point jump model.

The simplest possible jump model has only one possible jump amplitude x_0; i.e., $p(x) = \delta(x - x_0)$, using the Dirac delta function. Then (16.8) tells us that $\lambda^Q = \lambda \exp[(\gamma-1)x_0]$, and $q(x) = p(x)$. Only the jump frequency is changed by the risk adjustment. We mentioned above that the utility maximization provides an equation for the equilibrium equity return, but didn't show it. Here it is: $\alpha - r = (1-\gamma)\sigma^2 + (k\lambda - k^Q\lambda^Q)$. This is true for any model, but applied to the point jump case, it becomes

$$\begin{aligned}\alpha - r &= (1-\gamma)\sigma^2 + \lambda(\mathrm{e}^{x_0} - 1)(1 - \mathrm{e}^{(\gamma-1)x_0}) \\ &= (1-\gamma)(\sigma^2 + \lambda x_0^2) + O(x_0^3).\end{aligned}$$

Suppose the jumps are small, so the $O(x_0^3)$ terms are negligible. Then recall that the total statistical (actual) volatility is given by $\sigma_{total}^2 = \sigma^2 + \lambda(\mu_J^2 + \sigma_J^2) = \sigma^2 + \lambda x_0^2 + O(x_0^3)$. So, to leading order, the equity return relation is just $\alpha - r = (1-\gamma)\sigma_{total}^2$. But small jumps look a lot like Brownian motion anyway. So this is comforting to see that the classical relation $\alpha - r = (1-\gamma)\sigma^2$ is just being updated to include the extra jump volatility.

16.6 Exponential Lévy processes

From (16.10), $S_T = S_0 \exp\{\omega T + \sigma B(T) + \sum x_t\}$, where the sum ranges over all of the jumps that have occurred prior to T. In this last expression, ω and σ are constants, and the jumps are Poisson-driven described by the pair $[\lambda, p(x)]$ (actual) or $[\lambda^Q, q(x)]$ (risk-adjusted). This solution has the form $S_T = S_0 \exp X_T$, where X_T is a process of stationary, independent increments. Independent increments means that the random variables $X_{t(2)} - X_{t(1)}$ and $X_{t(4)} - X_{t(3)}$ are independent if $t(1) < t(2) < t(3) < t(4)$. Stationarity (or time-homogeneity) means that the probability distribution of $X_{t+s} - X_s$ does not depend upon s. One can ask: what is the most general process X_t that has stationary, independent increments?

The answer, which includes all our examples and an infinity of others, is the class of *Lévy* processes. All of the models we have discussed are examples of the stock price following an exponential (of a) Lévy process.

All Lévy processes belong to one of two types. (Warning: my category numbers don't agree with anybody's!) Type I consists of all of the linear combinations of Brownian motion plus Poisson-driven jumps. Type II is similar, but the jump process is not Poisson. There are too many small jumps; the Poisson intensity λ does not exist. So while λ^Q does not exist for Type II processes, you can take $\lambda^Q q(x) \to \mu(x)$ in many of the formulas. The non-negative $\mu(x)$ does exist and is called the (risk-adjusted) Lévy measure. Then, Type I means $\int \mu(x)dx < \infty$ and Type II means $\int \mu(x)dx = \infty$. Technically, there is a further sub-division of the Type II processes into (IIa, IIb), explained in the notes.[4] Our formulas are valid for Type I (the Poisson case) and Type IIa if you substitute $\lambda^Q q(x) \to \mu(x)$ everywhere.

While this is quite abstract, a good visualization for the exponential of a Type II process may well be the tick-by-tick changes of actively traded stocks. Consider Cisco or Intel or others that trade 25,000+ ticks per day, which is more than one trade a second during the normal U.S. session. So there are lots of small jumps. The many small jumps of Type II processes look a lot like Brownian motion watched under a strobe light. We're not going to say any more about Type II processes in this article, because we are more interested in the large "crash-type" jumps. These are well modeled by the Type I, Poisson-driven models.

[4] The Type IIa, IIb distinction hinges on why $\int \mu(x)dx = \infty$ in the first place. The integral always diverges for Type II processes because of the behavior of $\mu(x)$ near $x = 0$. A simple example is the Lévy α-stable process: $\mu(x) = c_\pm/|x|^{1+\alpha}$, $0 < \alpha < 2$, where $c_\pm$ are two constants for $x > 0$ and $x < 0$. Peter Carr and Liuren Wu [47] have proposed a special case of this model for stocks. In both this example and in general, the distinction between IIa and IIb is determined by the divergence as $x \to 0$. For the example, we can sub-divide $0 < \alpha < 2$, into $0 < \alpha < 1$ (IIa) and $1 \le \alpha < 2$ (IIb). The general rule is that $x\mu(x)$ can be integrated near the origin for Type IIa models. For Type IIb, $x\mu(x)$ is not integrable but $x^2\mu(x)$ is. See [163] for further discussion and references.

Table 16.1. Generalized Fourier Transforms for Various Financial Claims

Option	Payoff Function: $w(x)$	Payoff Transform: $\hat{w}(z)$	Strip of regularity: S_w
Call Option	$(\mathrm{e}^x - K)^+$	$-\frac{K^{1+\mathrm{i}z}}{z^2-\mathrm{i}z}$	$\Im z > 1$
Put Option	$(K - \mathrm{e}^x)^+$	$-\frac{K^{1+\mathrm{i}z}}{z^2-\mathrm{i}z}$	$\Im z < 0$
Covered Call	$\min(\mathrm{e}^x, K)$	$\frac{K^{1+\mathrm{i}z}}{z^2-\mathrm{i}z}$	$0 < \Im z < 1$
Delta Function	$\delta(x - \log K)$	$K^{\mathrm{i}z}$	Entire z-plane
Cash	1	$2\pi\delta(z)$	$\Im z = 0$

16.7 A formula for option values

In this section, we solve the integro-PDE (16.6). For the *particular* case, where $q(x)$ is a normal density, we already have Merton's solution at (16.4). For the particular case where $q(x)$ is a double exponential density, we refer the reader to Kou's solution. The trouble with these solutions is that, to find them, you need to rely upon special properties of their particular Lévy measures. Instead, here is an easier way to find the solution for every possible case at once. While we illustrate the method with the call option, it works just as well for any other payoff function. For more details, see [163].

Since $\mathrm{e}^{-rt}C(S,t)$ is a Q-martingale, we start with the well known expression

$$C(S_0, 0) = \mathrm{e}^{-rT}\mathbb{E}_0^Q\left[C(S_T, T)\right] = \mathrm{e}^{-rT}\mathbb{E}_0^Q\left[W(S_T)\right], \tag{16.12}$$

where $W(S_T) = \max(S_T - K, 0) = (S_T - K)^+$. The key idea is that everything that occurs in this problem has a simple *generalized* Fourier transform. Using the variable $x = \log(S_T)$, these transforms are $\hat{w}(z) = \int_{-\infty}^{\infty}\exp(\mathrm{i}zx)w(x)dx$, where $w(x) = W(\mathrm{e}^x)$ is the payoff function and z is a complex number. For example, since the call option payoff is $w(x) = (\mathrm{e}^x - K)^+$, then by a simple integration, $\hat{w}(z) = -K^{1+\mathrm{i}z}/(z^2 - \mathrm{i}z)$, $\Im z > 1$. Note that if z were real, this regular Fourier transforms would not exist. As shown in Lewis ([164]), the payoff transforms $\hat{w}(z)$ for typical claims exist and are regular in strips S_w in the complex z-plane. The transforms of many common payoffs are shown in Table 16.1.

Given $\hat{w}(z)$, the inversion formula is $w(x) = \int_{\mathrm{i}v-\infty}^{\mathrm{i}v+\infty}\exp(-\mathrm{i}zx)\hat{w}(z)dz/(2\pi)$. This is an integral along a line parallel to the real z-axis, where $z = u + \mathrm{i}v$. You just have to make sure that z stays in the strip of regularity. Again, for the call option, we found that we needed $v = \Im z > 1$.

So, the first step is to replace the payoff function in (16.12) by its transform:

$$\begin{aligned} C(S_0, 0) &= \frac{\mathrm{e}^{-rT}}{2\pi}\mathbb{E}_0^Q\left[\int_{\mathrm{i}v-\infty}^{\mathrm{i}v+\infty}\exp(-\mathrm{i}zx)\,\hat{w}(z)\,dz\right] \\ &= -\frac{\mathrm{e}^{-rT}}{2\pi}\mathbb{E}_0^Q\left[\int_{\mathrm{i}v-\infty}^{\mathrm{i}v+\infty}(S_T)^{-\mathrm{i}z}\frac{K^{1+\mathrm{i}z}}{z^2-\mathrm{i}z}\,dz\right], \quad (v > 1). \end{aligned} \tag{16.13}$$

The second step is to note that since $S_T = S_0\exp X_T$, where X_T is a Lévy process, we have $(S_T)^{-\mathrm{i}z} = (S_0)^{-\mathrm{i}z}\exp[-\mathrm{i}zX_T]$. Under a strip condition, explained in a footnote, it's permissible to bring the expectation symbol in (16.13) inside the integral. This gives us

$$C(S_0, 0) = -\frac{\mathrm{e}^{-rT}}{2\pi}\int_{\mathrm{i}v-\infty}^{\mathrm{i}v+\infty}(S_0)^{-\mathrm{i}z}\,\mathbb{E}_0^Q\left[\mathrm{e}^{-\mathrm{i}zX_T}\right]\frac{K^{1+\mathrm{i}z}}{z^2-\mathrm{i}z}\,dz. \tag{16.14}$$

Table 16.2. Some Characteristic Functions for Lévy Processes

Lévy Process	Characteristic Function: $\phi_T(z) = E[e^{izX_T}]$	Strip of regularity: S_X
Merton's Jump-diffusion	$\exp\left\{iz\omega T - \frac{1}{2}z^2\sigma^2 T + \lambda T\left(e^{iz\mu_J - z^2\sigma_J^2/2} - 1\right)\right\}$	Entire z-plane
Kou's Double Exponential	$\exp\left\{iz\omega T - \frac{1}{2}z^2\sigma^2 T + \lambda T\left(\frac{e^{iz\kappa}}{1+z^2\eta^2} - 1\right)\right\}$	$-\frac{1}{\eta} < \Im z < \frac{1}{\eta}$ $0 < \eta < 1$
Carr & Wu's Finite Moment Logstable	$\exp\{iz\omega T - (iz\sigma)^\alpha T \sec(\pi\alpha/2)\}$	$\Im z < 0$

For every Lévy process, $\phi_T(z) := \mathbb{E}_0^Q\left[e^{izX_T}\right]$ is called the characteristic function of the process. They are easy to compute for all of our examples, and typically have very short expressions, easily evaluated. Table 16.2 shows the results for the models discussed in this article. The parameters shown in the table are the Q-process parameters.

So the simple formula for the call option value, with z in a strip[5], is

$$C(S_0, 0) = -\frac{e^{-rT}}{2\pi}\int_{iv-\infty}^{iv+\infty} (S_0)^{-iz}\phi_T(-z)\,\frac{K^{1+iz}}{z^2 - iz}\,dz. \tag{16.15}$$

Notice how different this looks from, say (16.4), but it's identical for that model. For a general option, whose payoff function has the transform $\hat{w}(z)$, denote the Euro-style option value at $t = 0$ by $V(S_0)$. Then, the formula becomes

$$V(S_0) = \frac{e^{-rT}}{2\pi}\int_{iv-\infty}^{iv+\infty} (S_0)^{-iz}\phi_T(-z)\,\hat{w}(z)\,dz. \tag{16.16}$$

Similar or related results have been obtained by several authors, including Raible [215], Carr and Madan [46], Bakshi and Madan [14], and Boyarchenko and Levendorskii [35].

The integrals in (16.15) or (16.16) are straightforward and fast to compute. I have checked the formula numerically and found agreement with several known results: (i) the power series in solution of Merton's lognormal jump-diffusion given at (1.2); (ii) the special function solution for the double exponential model, using the Mathematica code that Steve Kou has made available on the Internet; (iii) the finite moment logstable model, using some high precision call option results graciously supplied by Liuren Wu and Peter Carr.

To conclude, I show a simple formula for the (risk-adjusted) characteristic functions, valid for any jump-diffusion:

$$\phi_T(z) = \exp\left\{T\left(iz\omega - \frac{1}{2}z^2\sigma^2 + \lambda^Q\int (e^{izx} - 1)q(x)\,dx\right)\right\}, \tag{16.17}$$

where $\omega = r - \sigma^2/2 - \lambda^Q \int (e^x - 1)\,q(x)\,dx$.

The integral in (16.17) is usually easy, and that's how the Table 16.2 entries were obtained. To derive (16.17), you just take the expectation of (16.11), using $p = iz$. For the details, see Lewis [163].

[5] Here's the strip rule. The functions $\phi_T(-z)$ exist in a strip easily deduced from the Table 16.2 entries. Just make sure this strip overlaps the payoff function strip from Table 16.1. You can integrate anywhere in the overlap.