Lorenzo Naranjo November 4, 2026

Pricing with Jumps

The geometric Brownian motion of the previous notes has continuous paths, so a stock price can never move by a large amount in a short time. Real prices do. Stocks gap on earnings announcements, markets crash, and firms default, and options markets price these risks: out-of-the-money puts trade at much higher implied volatilities than the Black-Scholes model can explain. Adding the jumps of the Poisson process to the Brownian motion of the stock price produces a *jump-diffusion*, the model of Merton (1976).

These notes bring together the two halves of the course. The Poisson process supplies the jumps, its exponential martingale supplies the change of measure, and its characteristic function supplies the pricing method. We first describe the Merton model and the distribution of its returns, extend Itô's formula to processes with jumps, and show how to change the intensity of a Poisson process. We then apply these tools to credit spreads and to options, where jumps make the market incomplete. The last section prices options by inverting the characteristic function of the log price, the method used in practice for models with jumps. These notes draw on Shreve (2004) and Duffie and Singleton (1999).

The Merton Model

Stock Prices with Jumps

Let $\{N_t\}$ be a Poisson process with intensity λ, and let $Y_1, Y_2, \ldots$ be iid $\mathcal{N}(m, s^2)$ random variables, independent of N and of the Brownian motion B. In the model of Merton (1976), the stock price is

$$S_t = S_0 \exp\left(\left(\mu - \frac{1}{2}\sigma^2 - \lambda k\right)t + \sigma B_t + \sum_{i=1}^{N_t} Y_i\right), \qquad k = \mathrm{E}(e^Y) - 1 = e^{m+\frac{1}{2}s^2} - 1. \quad (1)$$

Between jumps, the price is a geometric Brownian motion with volatility σ. At the time of the i-th event of the Poisson process, the log price jumps by Y_i, so the price is multiplied by e^{Y_i}, a percentage change of $e^{Y_i} - 1$. The constant k is the expected percentage jump, and with $m < 0$ jumps are mostly down, as in a crash.

The term $-\lambda k$ in the drift ensures that μ is still the expected return. Since B, N, and the Y_i are independent,

$$\mathrm{E}(S_t) = S_0 e^{(\mu-\lambda k)t}\,\mathrm{E}\left(e^{\sigma B_t - \frac{1}{2}\sigma^2 t}\right)\mathrm{E}\left(e^{\sum_{i=1}^{N_t} Y_i}\right) = S_0 e^{(\mu-\lambda k)t} \times 1 \times e^{\lambda t(M_Y(1)-1)} = S_0 e^{\mu t},$$

where the last expectation is the MGF of a compound Poisson process from the notes on Poisson processes, evaluated at $\theta = 1$, and $M_Y(1) = 1 + k$. Jumps add to the expected return at the rate λk, so without the correction, the stock would earn $\mu + \lambda k$. Applying the same argument conditionally on $\mathcal{F}_t$, the process $S_t e^{-\mu t}$ is a martingale.

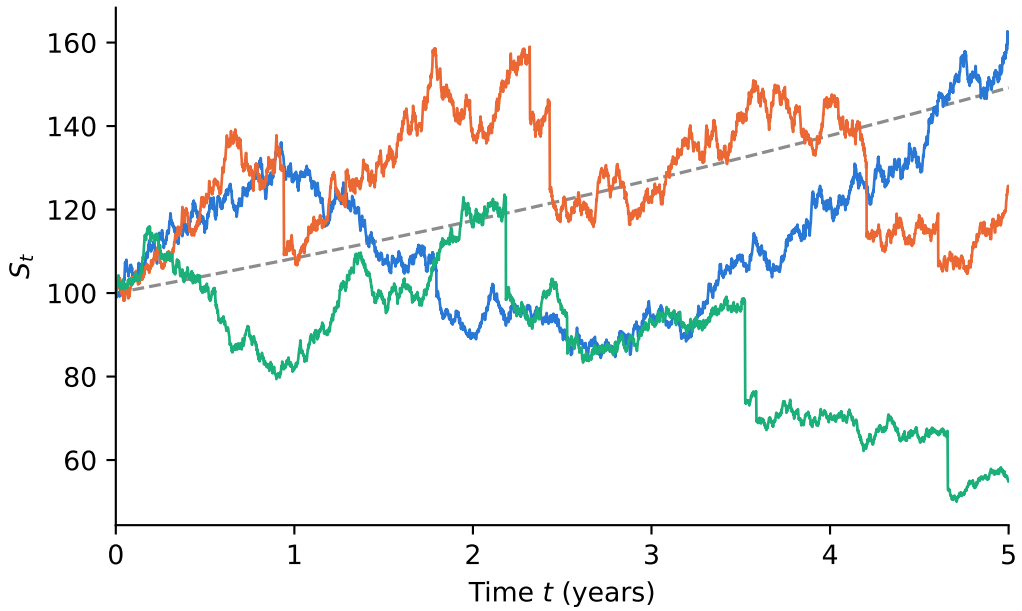


Figure 1: The figure plots three simulated paths of the Merton model with $S_0 = 100$, $\mu = 0.08$, $\sigma = 0.15$, $\lambda = 1$, $m = -0.08$, and $s = 0.10$ over five years. The dashed line is the expected price $S_0 e^{\mu t}$.

The Distribution of Returns

The log return $X_t = \ln(S_t/S_0)$ is the sum of a normal random variable and an independent compound Poisson process. Its characteristic function is the product of the two, using the example on the Merton model in the notes on Poisson processes:

$$\phi_{X_t}(u) = \exp\left(iu\left(\mu - \frac{1}{2}\sigma^2 - \lambda k\right)t - \frac{1}{2}u^2\sigma^2 t + \lambda t\left(e^{ium - \frac{1}{2}u^2 s^2} - 1\right)\right). \qquad (2)$$

Conditional on $N_t = n$, the log return is normal with mean $\left(\mu - \frac{1}{2}\sigma^2 - \lambda k\right)t + nm$ and variance $\sigma^2 t + ns^2$. Its density is therefore a mixture of normal densities, weighted by the Poisson probabilities of n jumps.

The cumulant generating function $K(\theta) = \ln \mathrm{E}(e^{\theta X_t})$ is the logarithm of the same expression with iu replaced by θ. The normal part contributes only to the first two cumulants, so every cumulant beyond the variance comes from the jumps. Writing κ_j for the j-th cumulant, the Poisson part contributes $\lambda t\ \mathrm{E}(Y^j)$ to each, so

$$\kappa_2 = \sigma^2 t + \lambda t(m^2 + s^2), \qquad \kappa_3 = \lambda t\left(m^3 + 3ms^2\right), \qquad \kappa_4 = \lambda t\left(m^4 + 6m^2 s^2 + 3s^4\right).$$

The skewness $\kappa_3/\kappa_2^{3/2}$ is negative when $m < 0$, and the excess kurtosis κ_4/κ_2^2 is positive, so the tails are heavier than normal. Both shrink as the horizon grows, like $1/\sqrt{t}$ and $1/t$. Over long horizons many jumps average out and returns look normal, but over short horizons a single jump dominates. This is why jumps matter most for short-dated options.

With the parameters of Figure 1, the one-month log return has skewness −1.31 and excess kurtosis 5.75. A normal distribution with the same variance underestimates the probability of large losses by orders of magnitude, as Figure 2 shows.

Itô's Formula with Jumps

Processes with Jumps

Consider a process that follows an Itô process between jumps and jumps by Z_i at the i-th event of a Poisson process:

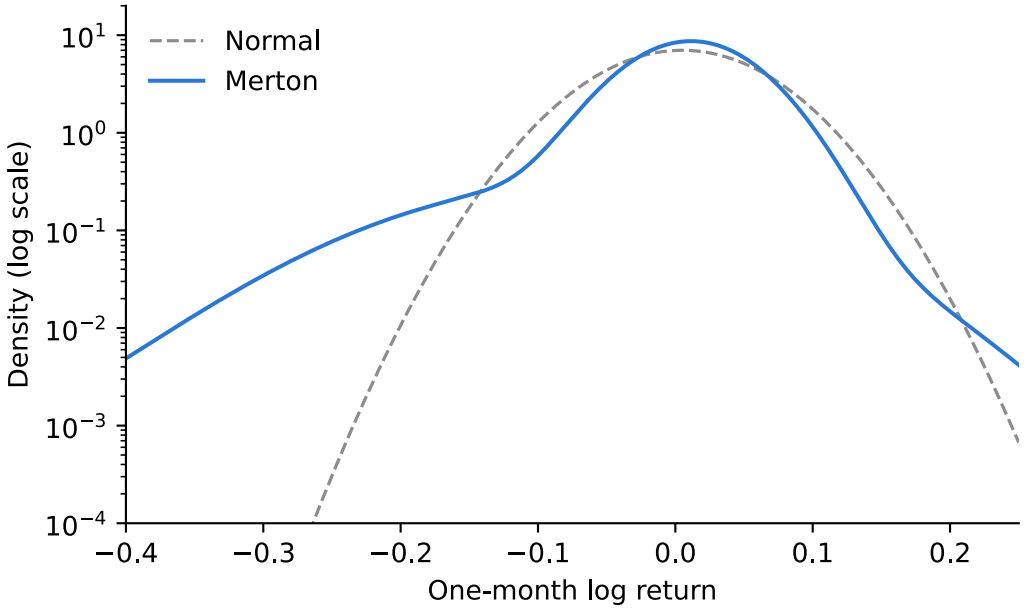


Figure 2: The figure plots the density of the one-month log return in the Merton model of Figure 1, together with a normal density with the same mean and variance, on a logarithmic scale.

$$X_t = X_0 + \int_0^t a_s \, \mathrm{d}s + \int_0^t b_s \, \mathrm{d}B_s + \sum_{i=1}^{N_t} Z_i.$$

Write X_{t-} for the value just before time t and $\Delta X_t = X_t - X_{t-}$ for the jump at time t, which is zero except at the event times. The log price in (1) has this form, with constant coefficients and $Z_i = Y_i$.

Over a short interval $\mathrm{d}t$, an event occurs with probability $\lambda \, \mathrm{d}t + o(\mathrm{d}t)$, so the expected change from the jumps is $\lambda \, \mathrm{E}(Z) \, \mathrm{d}t$. This gives the rule that replaces "the drift is zero" for processes with jumps.

Compensated Jumps

If $Z_1, Z_2, \ldots$ are iid with finite mean and independent of the Poisson process N with intensity λ, then

$$\sum_{i=1}^{N_t} Z_i - \lambda \, \mathrm{E}(Z) t$$

is a martingale. Consequently, a process with jumps is a martingale when its continuous drift plus λ times its expected jump is zero.

The proof is the same as for the compensated Poisson process $N_t - \lambda t$. The compound Poisson process has independent increments, and its expected increase over $(s, t]$ is $\lambda(t-s)\mathrm{E}(Z)$ by the moments of a compound Poisson process. The martingale $N_t - \lambda t$ is the case $Z_i = 1$. As in the notes on Poisson processes, the expected jump $\lambda \, \mathrm{E}(Z) \, \mathrm{d}t$ is called the *compensator* of the jumps.

The Formula

Between events, X is an ordinary Itô process, and Itô's formula from the notes on stochastic calculus applies. At an event time, X jumps from X_{t-} to $X_{t-} + \Delta X_t$, and a function of X jumps from $f(X_{t-})$ to $f(X_{t-} + \Delta X_t)$. There are finitely many events in any interval, so adding the two kinds of changes gives the formula.

Itô's Formula with Jumps

For a function $f(x)$ that is twice continuously differentiable,

$$f(X_t) - f(X_0) = \int_0^t \left(a_s f'(X_s) + \frac{1}{2} b_s^2 f''(X_s) \right) \mathrm{d}s + \int_0^t b_s f'(X_s) \, \mathrm{d}B_s$$
$$+ \sum_{s \le t} (f(X_{s-} + \Delta X_s) - f(X_{s-})),$$

where the last sum runs over the event times up to t.

The jump term uses the exact change in f, not a Taylor expansion, because a jump is not small. The quadratic variation of the jumps is $\sum(\Delta X_s)^2$, which was the result $[N,N]_t = N_t$ in the notes on Brownian motion, and the full change $f(X_{s-} + \Delta X_s) - f(X_{s-})$ accounts for this and all higher powers at once.

Example 1. Apply the formula to $S_t = e^{X_t}$, where X_t is the log price in (1), with $a = \mu - \frac{1}{2}\sigma^2 - \lambda k$, $b = \sigma$, and jumps Y_i. Between jumps, Itô's formula gives $\mathrm{d}S_t = S_t((\mu - \lambda k)\,\mathrm{d}t + \sigma\,\mathrm{d}B_t)$, as for a GBM. At a jump, S changes by $e^{X_{t-}+Y_i} - e^{X_{t-}} = S_{t-}(e^{Y_i} - 1)$. In differential notation,

$$\frac{\mathrm{d}S_t}{S_{t-}} = (\mu - \lambda k)\,\mathrm{d}t + \sigma\,\mathrm{d}B_t + \left(e^Y - 1\right)\mathrm{d}N_t,$$

where $\mathrm{d}N_t$ equals one at an event and zero otherwise. The expected percentage jump is k and jumps arrive at rate λ, so the compensator of the jump term is $\lambda k\,\mathrm{d}t$. The total expected return is $\mu - \lambda k + \lambda k = \mu$, as computed directly above. □

The same reasoning extends the product rule. At a common jump time of two processes, their product changes by $\Delta(XY) = X_{t-}\Delta Y + Y_{t-}\Delta X + \Delta X\,\Delta Y$, which is the exact change of xy and has the same form as the continuous product rule.

Changing the Jump Intensity

Girsanov's theorem for Brownian motion reweights paths with the exponential martingale $e^{-\theta B_t - \frac{1}{2}\theta^2 t}$ and shifts the drift. The Poisson process has its own exponential martingale, $e^{\theta N_t - \lambda t(e^\theta - 1)}$, and reweighting paths with it changes the intensity.

Changing the Intensity

Let N be a Poisson process with intensity λ, fix $\lambda^* > 0$ and $T > 0$, and let

$$Z_t = \left(\frac{\lambda^*}{\lambda}\right)^{N_t} e^{-(\lambda^* - \lambda)t}, \qquad 0 \le t \le T.$$

Then Z is a positive martingale with $Z_0 = 1$, and under the measure $\mathrm{P}^*(A) = \mathrm{E}(Z_T \mathbf{1}_{\{A\}})$, the process N is a Poisson process with intensity λ^*.

Proof

With $e^{\theta} = \lambda^*/\lambda$, we have $\lambda(e^{\theta} - 1) = \lambda^* - \lambda$, so Z_t is the exponential martingale of the Poisson process. As in the proof of Girsanov's theorem, $\mathrm{E}^*(W) = \mathrm{E}(Z_t W)$ for any W known at time t. For $s < t$ and real u and v, write $\Delta N = N_t - N_s$ and split Z_t into a factor depending on N_s and a factor depending on ΔN, which are independent. For the first factor,

$$\mathrm{E}\left(\left(\frac{\lambda^*}{\lambda}\right)^{N_s} e^{-(\lambda^*-\lambda)s} e^{uN_s}\right) = e^{-(\lambda^*-\lambda)s}\, e^{\lambda s\left(\frac{\lambda^*}{\lambda}e^{u}-1\right)} = e^{\lambda^* s(e^{u}-1)},$$

using the MGF of a Poisson random variable. The second factor equals $e^{\lambda^*(t-s)(e^{v}-1)}$ in the same way. Therefore, the joint MGF of N_s and ΔN under P^* factors into the MGFs of independent Poisson random variables with means $\lambda^* s$ and $\lambda^*(t-s)$, which is the definition of a Poisson process with intensity λ^*. The same calculation applies to any finite number of increments. □

A Brownian change of measure can only shift the drift, because the quadratic variation of Brownian motion is the same on every path. The quadratic variation of the Poisson process, $[N, N]_t = N_t$, is random, so reweighting paths can change how often jumps occur. Paths with many events receive more weight when $\lambda^* > \lambda$, so jumps become more frequent. Since $Z_T > 0$, the two measures agree on which events are impossible.

The same idea can also change the distribution of the jump sizes. If the sizes have density g under P and we want density g^* under P^*, we multiply Z_t by $\prod_{i \le N_t} g^*(Y_i)/g(Y_i)$. Brownian motion and jumps can be reweighted at the same time, since independent sources of randomness can be reweighted separately. In general, a change of measure in a jump-diffusion can shift the Brownian drift, change the intensity, and change the distribution of jump sizes.

Application: Credit Spreads

The notes on Poisson processes priced a defaultable zero-coupon bond by discounting its expected payoff at the risk-free rate, obtaining a credit spread equal to the default intensity. We can now include a premium for default risk.

A firm defaults at the first event τ of a Poisson process with intensity λ under P. Defaults tend to happen in bad times, when a dollar is more valuable, so we let the SDF jump up at default. Specifically, let

$$\Lambda_t = e^{-rt}(1+\eta)^{N_t} e^{-\lambda\eta t}, \qquad \eta > -1,$$

which jumps by the factor $1+\eta$ at each event. The second and third factors form the exponential martingale of the Poisson process with $e^{\theta} = 1+\eta$, so $\Lambda_t e^{rt}$ is a martingale and the SDF has drift $-r$, as required. When $\eta > 0$, a dollar received just after default is worth more than a dollar received just before.

The bond pays one dollar at T if the firm survives. At default, the bond loses a fraction L of its value, so bondholders recover the fraction $1-L$, and we suppose the recovered amount is then invested at the risk-free rate. Guess that before default the bond price is $P_t = e^{-c(T-t)}$ for some yield c. The pricing equation of the notes on discount factors requires ΛP to have zero expected change. Before default, ΛP has continuous drift

$$\Lambda P(-r - \lambda\eta + c)\,\mathrm{d}t.$$

At default, Λ is multiplied by $1+\eta$ and P by $1-L$, so ΛP changes by $\Lambda P((1+\eta)(1-L)-1)$, with intensity λ. Setting the total expected change to zero,

$$-r - \lambda\eta + c + \lambda((1+\eta)(1-L)-1) = 0 \quad \Leftrightarrow \quad c = r + \lambda(1+\eta)L.$$

Writing $\lambda^* = \lambda(1+\eta)$, the credit spread is

$$c - r = \lambda^* L. \tag{3}$$

The same result follows from a change of measure. The risk-neutral density is $\Lambda_t\beta_t = (1+\eta)^{N_t}e^{-\lambda\eta t}$, which is the density Z_t of the previous section with $\lambda^*/\lambda = 1+\eta$. Under the risk-neutral measure, default arrives with intensity $\lambda^* = \lambda(1+\eta)$, and the bond is priced by discounting at the risk-free rate. With zero recovery, $L = 1$, this gives $P = e^{-(r+\lambda^*)T}$, the formula of the notes on Poisson processes with the risk-neutral intensity λ^* in place of the historical intensity λ.

The spread (3) exceeds the expected loss rate λL by the factor $1+\eta$, the compensation investors require for bearing default risk that is concentrated in bad times. Spreads on

investment-grade corporate bonds are typically several times the expected loss rate, which points to large values of η.

Example 2. A firm defaults with historical intensity $\lambda = 0.01$ per year, and bondholders lose $L = 0.6$ of the bond's value at default. The expected loss rate is $\lambda L = 0.006$, or 60 basis points per year. If the SDF jumps by the factor $1 + \eta = 3$ at default, the risk-neutral intensity is $\lambda^* = 0.03$, and the credit spread is

$$\lambda^* L = 0.03 \times 0.6 = 0.018,$$

or 180 basis points, three times the expected loss. □

Option Pricing in the Merton Model

Jumps Make the Market Incomplete

In the Black-Scholes model, an option can be replicated by trading the stock, because the option and the stock respond to the same single source of risk. With jumps this fails. A position in the stock can offset the exposure of an option to $\mathrm{d}B$, but not its exposure to a jump of random size: the change in an option's value at a jump is a nonlinear function of the jump size, while the change in the stock position is linear in $e^Y - 1$. The market is *incomplete*, and no-arbitrage alone does not determine option prices.

In terms of the previous section, many SDFs price the stock correctly. Each one implies a risk-neutral measure under which the discounted stock price is a martingale, which by the compensated jump rule means

$$\frac{\mathrm{d}S_t}{S_{t-}} = (r - \lambda^* k^*)\,\mathrm{d}t + \sigma\,\mathrm{d}B_t^* + \left(e^Y - 1\right)\mathrm{d}N_t,$$

where under P^* the jumps arrive at intensity λ^* and have expected percentage size k^*. Different choices of λ^* and of the jump distribution give different option prices. Choosing among them is an economic assumption about how investors price jump risk, or, in practice, a matter of calibrating the risk-neutral parameters to observed option prices.

Merton (1976) assumed that jump risk is specific to the firm and can be diversified away, so investors require no premium for it. The SDF then does not jump, and only

the Brownian drift changes under the risk-neutral measure. The intensity and the jump distribution are the same under P and P^*, and

$$\frac{\mathrm{d}S_t}{S_{t-}} = (r - \lambda k)\,\mathrm{d}t + \sigma\,\mathrm{d}B_t^* + \left(e^Y - 1\right)\mathrm{d}N_t.$$

Under this measure, S_T has the form (1) with μ replaced by r.

Merton's Formula

Conditional on $N_T = n$, the log price under P^* is normal with mean $\ln S_0 + \left(r - \lambda k - \frac{1}{2}\sigma^2\right)T + nm$ and variance $\sigma^2 T + ns^2$. A call is therefore a Poisson mixture of options in Black-Scholes models with different volatilities and interest rates.

Merton's Option Pricing Formula

Let $C^{BS}(S, K, r, \sigma, T)$ denote the Black-Scholes call price. In the Merton model with unpriced jump risk, the call price is

$$C = \sum_{n=0}^{\infty} e^{-\lambda' T}\frac{(\lambda' T)^n}{n!}\,C^{BS}(S_0, K, r_n, \sigma_n, T),$$

where $\lambda' = \lambda(1 + k)$, $\sigma_n^2 = \sigma^2 + ns^2/T$, and $r_n = r - \lambda k + n\ln(1 + k)/T$.

Proof

In a Black-Scholes model with rate r_n and volatility σ_n, the log price is normal with mean $\ln S_0 + (r_n - \frac{1}{2}\sigma_n^2)T$ and variance $\sigma_n^2 T$. Since $\ln(1 + k) = m + \frac{1}{2}s^2$,

$$\left(r_n - \frac{1}{2}\sigma_n^2\right)T = (r - \lambda k)T + n\left(m + \frac{1}{2}s^2\right) - \frac{1}{2}\sigma^2 T - \frac{1}{2}ns^2 = \left(r - \lambda k - \frac{1}{2}\sigma^2\right)T + nm,$$

and $\sigma_n^2 T = \sigma^2 T + ns^2$. Conditional on $N_T = n$, the log price therefore has the same distribution as in that Black-Scholes model. The Black-Scholes price discounts at r_n, so

$$e^{-rT}\mathrm{E}^*((S_T - K)^+ \mid N_T = n) = e^{(r_n - r)T}\,C^{BS}(S_0, K, r_n, \sigma_n, T).$$

Averaging over the Poisson distribution of N_T,

$$C = \sum_{n=0}^{\infty} e^{-\lambda T} \frac{(\lambda T)^n}{n!} e^{-\lambda k T} (1+k)^n \, C^{BS}(S_0, K, r_n, \sigma_n, T),$$

and $e^{-\lambda T} e^{-\lambda k T} (\lambda T)^n (1+k)^n = e^{-\lambda' T} (\lambda' T)^n$. □

Example 3. Take the Merton model of Figure 1 with $S_0 = 100$, $r = 0.05$, and a three-month at-the-money call, $K = 100$ and $T = 0.25$. Merton's formula gives a price of 4.3499, compared with 3.6351 in a Black-Scholes model with the diffusion volatility $\sigma = 0.15$ alone. Inverting the Black-Scholes formula, the Merton price corresponds to an implied volatility of 0.1865. A three-month put with strike 80 is worth 0.1342 in the Merton model and only 0.0015 without jumps. □

The implied volatility depends on the strike. Figure 3 plots it for three-month options. Low strikes, which pay off after a crash, have the highest implied volatilities, because downward jumps make large losses far more likely than the lognormal distribution allows. This pattern, called the volatility *skew*, has been a persistent feature of index options since the crash of 1987, and the Black-Scholes model, with a single volatility for all strikes, cannot produce it.

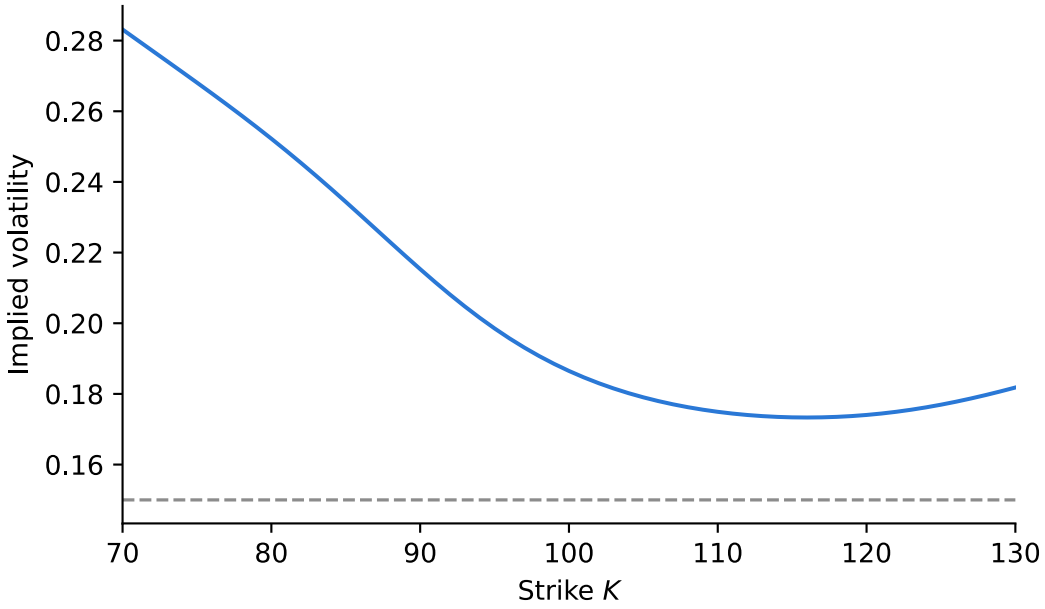


Figure 3: The figure plots the Black-Scholes implied volatility of three-month Merton call prices against the strike, with $S_0 = 100$, $r = 0.05$, and the jump parameters of Figure 1. The dashed line is the diffusion volatility $\sigma = 0.15$.

Fourier Pricing

Prices from Characteristic Functions

Merton's formula works because, given the number of jumps, the log price is normal. For most models with jumps or stochastic volatility the density of the log price is not known in closed form, but its characteristic function is, as (2) illustrates. The notes on generating functions noted that probabilities can be recovered from the characteristic function by inversion. This is the basis of option pricing in such models.

The pricing notes wrote the price of a call as

$$C = S_0 \mathrm{P}^S(S_T > K) - Ke^{-rT}\mathrm{P}^*(S_T > K),$$

using the risk-neutral measure P^* and the measure P^S that uses the stock as the unit of account. The derivation did not depend on the Black-Scholes model, so the formula holds in any model with a risk-neutral measure. Both probabilities can be computed from the characteristic function of $X = \ln S_T$ under P^*.

Inversion Formula

If X has characteristic function ϕ and a continuous distribution, then for every real x

$$\mathrm{P}(X > x) = \frac{1}{2} + \frac{1}{\pi}\int_0^\infty \frac{\mathrm{Im}\left(e^{-iux}\phi(u)\right)}{u}\,\mathrm{d}u.$$

Proof

The Dirichlet integral $\int_0^\infty \frac{\sin(au)}{u}\,\mathrm{d}u$ equals $\frac{\pi}{2}$ if $a > 0$ and $-\frac{\pi}{2}$ if $a < 0$. Therefore, for $y \neq x$,

$$\mathbf{1}_{\{y>x\}} = \frac{1}{2} + \frac{1}{\pi}\int_0^\infty \frac{\sin(u(y-x))}{u}\,\mathrm{d}u.$$

Setting $y = X$, taking expectations, and interchanging the expectation and the integral, which can be justified,

$$\mathrm{P}(X > x) = \frac{1}{2} + \frac{1}{\pi}\int_0^\infty \frac{\mathrm{E}\left(\sin(u(X-x))\right)}{u}\,\mathrm{d}u.$$

Finally, $\mathrm{E}(\sin(u(X-x)))$ is the imaginary part of $\mathrm{E}(e^{iu(X-x)}) = e^{-iux}\phi(u)$. □

The formula gives $\mathrm{P}^*(S_T > K)$ directly, with $x = \ln K$ and ϕ the risk-neutral characteristic function of $\ln S_T$. For P^S we need the characteristic function of $\ln S_T$ under the stock measure, whose density with respect to P^* is $S_T/\mathrm{E}^*(S_T) = e^X/\phi(-i)$, with $\phi(-i) = \mathrm{E}^*(e^X) = S_0 e^{rT}$. Therefore,

$$\phi^S(u) = \mathrm{E}^*\left(\frac{e^X}{\phi(-i)}e^{iuX}\right) = \frac{\phi(u-i)}{\phi(-i)}.$$

The change of numeraire, which in the Black-Scholes model shifted the drift of the log price by σ^2, becomes a shift in the argument of the characteristic function.

Fourier Pricing Formula

Let ϕ be the characteristic function of $\ln S_T$ under P^*. Then $C = S_0\Pi_1 - Ke^{-rT}\Pi_2$, where

$$\Pi_1 = \frac{1}{2} + \frac{1}{\pi}\int_0^\infty \frac{\mathrm{Im}\left(e^{-iu\ln K}\phi(u-i)\right)}{u\,\phi(-i)}\,\mathrm{d}u, \qquad \Pi_2 = \frac{1}{2} + \frac{1}{\pi}\int_0^\infty \frac{\mathrm{Im}\left(e^{-iu\ln K}\phi(u)\right)}{u}\,\mathrm{d}u.$$

Example 4. For the at-the-money call of Example 3, numerical integration of the Fourier pricing formula with the characteristic function (2), with μ replaced by r, gives $\Pi_1 = 0.6092$ and $\Pi_2 = 0.5728$, and a call price of

$$C = 100 \times 0.6092 - 100e^{-0.05\times 0.25} \times 0.5728 = 4.3499,$$

which matches the value 4.3499 from Merton's formula. □

Why Fourier Methods

The Fourier formula needs nothing from the model except the characteristic function of the log price. Any model in which that function has a closed form can be priced the same way, including models with other jump distributions and the stochastic volatility model of Heston (1993). This is why the characteristic function, which in the notes on Poisson processes was a convenient way to describe a compound Poisson process, becomes the central object in modern option pricing. In practice, related formulas evaluated with

the fast Fourier transform price options at many strikes at once, which makes it fast to calibrate a model's risk-neutral parameters to an entire set of option prices.

References

Duffie, Darrell, and Kenneth J. Singleton. 1999. "Modeling Term Structures of Defaultable Bonds." *Review of Financial Studies* 12 (4): 687–720.

Heston, Steven L. 1993. "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options." *Review of Financial Studies* 6 (2): 327–43.

Merton, Robert C. 1976. "Option Pricing When Underlying Stock Returns Are Discontinuous." *Journal of Financial Economics* 3 (1–2): 125–44.

Shreve, Steven E. 2004. *Stochastic Calculus for Finance II: Continuous-Time Models*. Springer Finance. Springer.