Lorenzo Naranjo November 6, 2026

Poisson Processes

Many events in finance arrive at unpredictable moments: a firm defaults, an insurance claim is filed, an order reaches the exchange, a central bank surprises the market, or a stock price jumps on news. The *Poisson process* is the basic model for the times at which such events occur. It counts how many events have happened up to each point in time, and it rests on two assumptions: events arrive at a constant average rate, and what happens over one period of time tells us nothing about what will happen over another.

The Poisson process is the first process in continuous time that we study. It arises from the Poisson limit of sums of Bernoulli random variables, its paths are simple to draw, and most calculations require nothing beyond generating functions and martingales. Its continuous counterpart, a process that moves all the time by tiny amounts instead of by jumps, is Brownian motion.

These notes draw on Grimmett and Stirzaker (2001) and on the treatment of jump processes in Shreve (2004).

Counting Events in Continuous Time

A Bernoulli Construction

Suppose that events can occur at any moment, at an average rate of λ events per unit of time. To model this, divide the interval $[0, t]$ into n short slots of length $\Delta t = t/n$. In each slot, an event occurs with probability $\lambda \Delta t$ and no event occurs with probability $1 - \lambda \Delta t$, independently across slots. Let X_j be the Bernoulli indicator that an event occurs in slot j. The number of events up to time t is

$$N_t^{(n)} = X_1 + X_2 + \cdots + X_n.$$

This is a binomial random variable with n trials and success probability $\lambda t/n$, so the expected number of events is λt regardless of n.

As the slots become shorter, the probability of an event in each slot shrinks while the number of slots grows. This is exactly the setting of the Poisson limit in the notes on generating functions. The PGF of $N_t^{(n)}$ is

$$G_{N_t^{(n)}}(z) = \left(1 + \frac{\lambda t(z-1)}{n}\right)^n \xrightarrow[n\to\infty]{} e^{\lambda t(z-1)},$$

so in the limit the number of events up to time t has a Poisson distribution with mean λt.

The construction gives us more than the distribution at a single date. For $s < t$, the events in $(s, t]$ are determined by the slots in that interval, which are independent of the slots in $[0, s]$. The number of events in $(s, t]$ is again binomial, and in the limit it is Poisson with mean $\lambda(t - s)$, a quantity that depends only on the length of the interval. These properties survive the limit and define the Poisson process.

Definition

Poisson Process

A stochastic process $\{N_t : t \geq 0\}$ is called a *Poisson process* with rate (or intensity) $\lambda > 0$ if it has the following three properties:

i. $N_0 = 0$.

ii. For any finite set of times $0 \leq t_0 < t_1 < \cdots < t_n$, the increments $N_{t_1} - N_{t_0}, N_{t_2} - N_{t_1}, \ldots, N_{t_n} - N_{t_{n-1}}$ are independent.

iii. For any $0 \leq s < t$, the increment $N_t - N_s$ has a Poisson distribution with mean $\lambda(t - s)$.

The paths of N_t are step functions that start at zero and jump up by one at each event. By convention the paths are right-continuous, so N_t includes an event that occurs exactly at time t.

Using the moments of the Poisson distribution, we have

$$\mathrm{E}(N_t) = \lambda t, \qquad \mathrm{V}(N_t) = \lambda t, \qquad G_{N_t}(z) = e^{\lambda t(z-1)}.$$

The rate λ is the expected number of events per unit of time. An insurer whose claims arrive at the rate $\lambda = 10$ per year, for example, expects ten claims a year.

The Bernoulli construction also describes the behavior of the process over a short interval. Over an interval of length $\mathrm{d}t$, the Poisson probabilities give

$$\begin{aligned}
\mathrm{P}(N_{t+\mathrm{d}t} - N_t = 0) &= e^{-\lambda\,\mathrm{d}t} = 1 - \lambda\,\mathrm{d}t + o(\mathrm{d}t), \\
\mathrm{P}(N_{t+\mathrm{d}t} - N_t = 1) &= \lambda\,\mathrm{d}t\,e^{-\lambda\,\mathrm{d}t} = \lambda\,\mathrm{d}t + o(\mathrm{d}t), \\
\mathrm{P}(N_{t+\mathrm{d}t} - N_t \geq 2) &= o(\mathrm{d}t),
\end{aligned}$$

where $o(\mathrm{d}t)$ denotes terms that vanish faster than $\mathrm{d}t$. Over a short interval there is either no event or a single event, and the probability of an event is proportional to the length of the interval. Writing $\mathrm{d}N_t = N_{t+\mathrm{d}t} - N_t$, this says that $\mathrm{d}N_t$ is approximately a Bernoulli random variable with $\mathrm{E}(\mathrm{d}N_t) = \lambda\,\mathrm{d}t$.

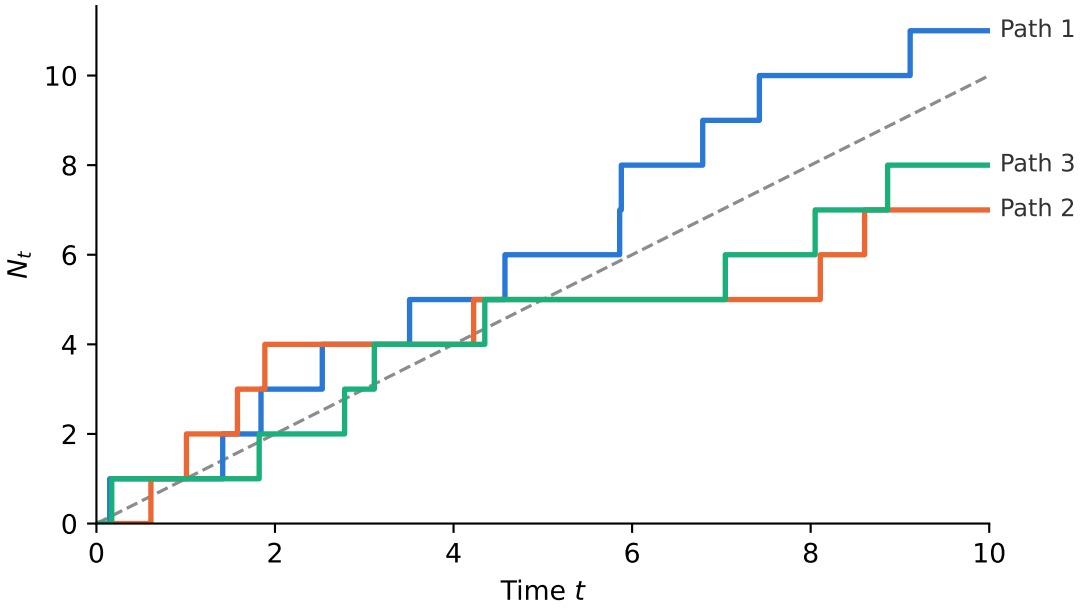


Figure 1: The figure plots three simulated paths of a Poisson process with rate $\lambda = 1$ over $0 \leq t \leq 10$. The dashed line denotes $\mathrm{E}(N_t) = \lambda t$.

Combining and Splitting Poisson Processes

Two operations on Poisson processes appear constantly in applications: merging several independent streams of events into one, and splitting one stream into several types. Both follow from the generating function results for Poisson sums and thinning.

Superposition and Thinning

i. **Superposition.** If $N_t^{(1)}$ and $N_t^{(2)}$ are independent Poisson processes with rates λ_1 and λ_2, then $N_t = N_t^{(1)} + N_t^{(2)}$ is a Poisson process with rate $\lambda_1 + \lambda_2$.

ii. **Thinning.** If each event of a Poisson process with rate λ is independently classified as type A with probability p and as type B with probability $1 - p$, then the counts N_t^A and N_t^B of each type are *independent* Poisson processes with rates λp and $\lambda(1 - p)$.

Figure 2 shows the classification behind thinning. Each event of the original process takes one of two branches, independently of the other events and of the times at which they arrive.

Proof

For superposition, the sum of independent processes with independent increments has independent increments, and by the notes on generating functions the increment over $(s, t]$ is Poisson with mean $(\lambda_1 + \lambda_2)(t - s)$.

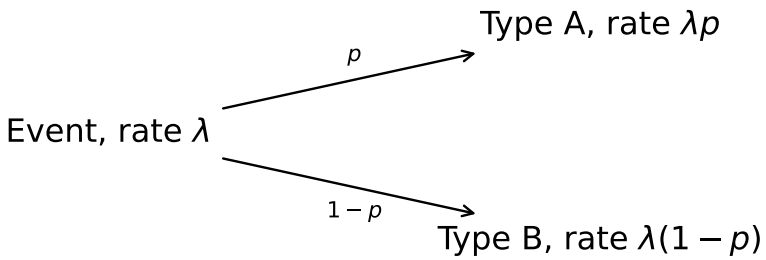


Figure 2: The figure shows how thinning classifies each event. An event of the process with rate λ is of type A with probability p and of type B with probability $1 - p$, which gives two processes with rates λp and $\lambda(1 - p)$.

For thinning, it is enough to consider a single interval $[0, t]$, since increments over disjoint intervals are built from disjoint sets of events. We compute the *joint* PGF of N_t^A and N_t^B by conditioning on N_t. Given $N_t = n$, each of the n events contributes a factor z_1 if it is of type A and z_2 if it is of type B, so

$$\mathrm{E}\left(z_1^{N_t^A} z_2^{N_t^B} \mid N_t = n\right) = (pz_1 + (1-p)z_2)^n.$$

Taking expectations over N_t and using its PGF,

$$\mathrm{E}\left(z_1^{N_t^A} z_2^{N_t^B}\right) = e^{\lambda t(pz_1+(1-p)z_2-1)} = e^{\lambda pt(z_1-1)}\, e^{\lambda(1-p)t(z_2-1)}.$$

The joint PGF factors into the PGF of a Poisson(λpt) random variable in z_1 times the PGF of a Poisson$(\lambda(1-p)t)$ random variable in z_2. Since a joint PGF determines the joint distribution, the two counts are independent Poisson random variables. □

The independence of the thinned processes is surprising at first. If we learn that many type A events occurred, we might expect fewer type B events. That intuition is correct when the total number of events is fixed, but here the total is random. Many type A events suggest that the total was large, which exactly offsets the smaller share left for type B. A practice problem in the notes on generating functions works through the same calculation for buy and sell orders arriving at a market maker, and compares it with a fixed number of orders.

Example 1. Suppose large news announcements about a stock arrive as a Poisson process with rate $\lambda = 6$ per year, and each announcement is good news with probability $p = 0.4$ and bad news otherwise. Good-news jumps then arrive at rate 2.4 per year and bad-news jumps at rate 3.6 per year, and the two processes are independent. The probability of a year with no bad-news jumps is $e^{-3.6} = 0.0273$, whatever the number of good-news jumps. □

Waiting Times

The First Event

Let T_1 be the time of the first event. The first event occurs after time t if and only if no event occurs in $[0, t]$, so

$$\mathrm{P}(T_1 > t) = \mathrm{P}(N_t = 0) = e^{-\lambda t}.$$

Therefore, T_1 has an exponential distribution with rate λ, and the expected waiting time for the first event is $1/\lambda$. A firm that defaults at the first event of a Poisson process with intensity $\lambda = 0.02$ per year waits on average $1/0.02 = 50$ years for its default.

The Bernoulli construction gives the same answer from a different angle. Each slot either holds the first event, with probability $\lambda\Delta t$, or passes without one, as in Figure 3. The number of slots until the first event is therefore geometric with success probability $\lambda\Delta t$, and the probability that no event occurs in the first $t/\Delta t$ slots, the probability of reaching row $t/\Delta t$ of the left column, is

$$(1 - \lambda\Delta t)^{t/\Delta t} \xrightarrow[\Delta t \to 0]{} e^{-\lambda t}.$$

The exponential distribution is the continuous-time limit of the geometric distribution, in the same way that the Poisson distribution is the limit of the binomial.

Memorylessness

The exponential distribution has a property that makes it the natural model for waiting times in the Poisson process. If $T \sim \mathrm{Exp}(\lambda)$, then for all $s, t \geq 0$

$$\mathrm{P}(T > s + t \mid T > s) = \frac{\mathrm{P}(T > s + t)}{\mathrm{P}(T > s)} = \frac{e^{-\lambda(s+t)}}{e^{-\lambda s}} = e^{-\lambda t} = \mathrm{P}(T > t).$$

Having waited s units of time without an event does not change the distribution of the remaining waiting time. This *memoryless property* characterizes the exponential distribution among continuous distributions on $[0, \infty)$, just as the geometric distribution is the only memoryless distribution on $\{1, 2, \ldots\}$.

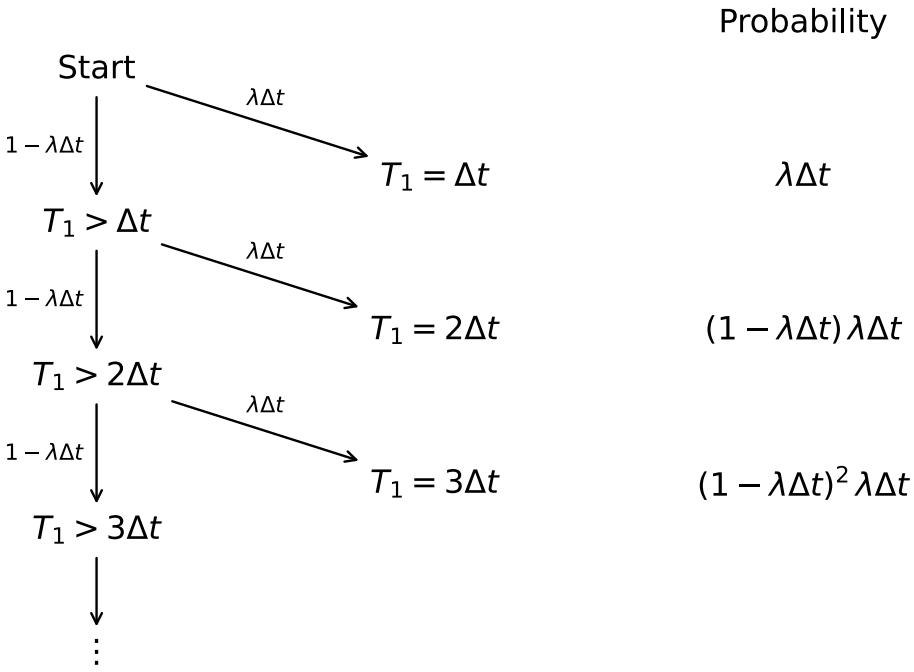


Figure 3: The figure shows the first event in the Bernoulli construction. In each slot of length Δt, the first event occurs with probability $\lambda\Delta t$, and with probability $1-\lambda\Delta t$ the process moves down the left column to the next slot. The right column gives the probability that the first event falls in each slot.

In credit risk, memorylessness means that the probability that a firm defaults over the next year does not depend on how long the firm has survived. The probability of default over a short interval, given survival up to time t, is

$$\mathrm{P}(T \le t + \mathrm{d}t \mid T > t) = 1 - e^{-\lambda\,\mathrm{d}t} = \lambda\,\mathrm{d}t + o(\mathrm{d}t),$$

which is why λ is also called the *hazard rate* of default. Models used in practice let the hazard rate vary over time and with the state of the economy, but the constant-intensity case is the building block.

Interarrival Times

Memorylessness is what lets the process start afresh after each event. Because increments are independent and stationary, once an event has occurred the past carries no information about the future, and the time until the next event is again exponential with rate λ, independent of everything that happened before. We state the result without a formal proof.

Interarrival Times

Let $T_1, T_2, \ldots$ denote the times between consecutive events of a Poisson process with rate λ. Then $T_1, T_2, \ldots$ are independent random variables, each with an $\text{Exp}(\lambda)$ distribution. Conversely, if $T_1, T_2, \ldots$ are iid $\text{Exp}(\lambda)$ and $W_n = T_1 + \cdots + T_n$, then

$$N_t = \max\{n \geq 0 : W_n \leq t\}$$

is a Poisson process with rate λ.

The converse gives a simple recipe to simulate a Poisson process: draw exponential waiting times and add them up. The paths in Figure 1 were generated in this way.

Arrival Times

The time of the n-th event is $W_n = T_1 + \cdots + T_n$, a sum of n independent exponential random variables. By the results in the notes on generating functions, W_n has a gamma distribution with

$$M_{W_n}(\theta) = \left(\frac{\lambda}{\lambda - \theta}\right)^n, \qquad \text{E}(W_n) = \frac{n}{\lambda}, \qquad \text{V}(W_n) = \frac{n}{\lambda^2}.$$

The arrival times and the counts are two descriptions of the same process. The n-th event has occurred by time t if and only if at least n events have occurred in $[0, t]$, so

$$\text{P}(W_n \leq t) = \text{P}(N_t \geq n) = 1 - \sum_{k=0}^{n-1} e^{-\lambda t} \frac{(\lambda t)^k}{k!}.$$

Example 2. A portfolio contains bonds issued by 50 firms, each defaulting at the first event of an independent Poisson process with intensity 0.02 per year. Until the first default occurs, defaults in the portfolio arrive as a Poisson process with rate $50 \times 0.02 = 1$ per year, by superposition. Ignoring the small reduction in the rate after each default, the probability that at least two firms default within a year is

$$\mathrm{P}(W_2 \le 1) = 1 - e^{-1}(1+1) = 1 - 2e^{-1} = 0.2642.$$

□

Martingales of the Poisson Process

Let $\mathcal{F}_t$ denote the information generated by observing the process $\{N_s : s \le t\}$ up to time t. The definition of a martingale carries over from discrete time: a process $\{M_t\}$ adapted to $\{\mathcal{F}_t\}$ is a martingale if $\mathrm{E}(M_t \mid \mathcal{F}_s) = M_s$ for all $s \le t$. Since increments of the Poisson process are independent of the past, conditional expectations of functions of $N_t - N_s$ given $\mathcal{F}_s$ reduce to ordinary expectations, exactly as for the random walk.

The Poisson process itself is not a martingale, since it only moves up. Its expected increase over $(s, t]$ is $\lambda(t-s)$, and subtracting it produces a martingale.

Martingales of the Poisson Process

Let $\{N_t\}$ be a Poisson process with rate λ. The following processes are martingales with respect to $\{\mathcal{F}_t\}$:

i. The *compensated Poisson process* $M_t = N_t - \lambda t$.

ii. The process $M_t^2 - \lambda t$.

iii. For any real θ, the exponential martingale $Z_t = \exp\big(\theta N_t - \lambda t(e^\theta - 1)\big)$.

Proof

Fix $s < t$ and write $\Delta N = N_t - N_s$, which is independent of $\mathcal{F}_s$ and has a Poisson distribution with mean $\lambda(t-s)$. Let $\Delta M = M_t - M_s = \Delta N - \lambda(t-s)$, so that $\mathrm{E}(\Delta M) = 0$ and $\mathrm{E}(\Delta M^2) = \mathrm{V}(\Delta N) = \lambda(t-s)$.

For the first process,

$$\mathrm{E}(M_t \mid \mathcal{F}_s) = M_s + \mathrm{E}(\Delta M) = M_s.$$

For the second, expanding $M_t^2 = M_s^2 + 2M_s\Delta M + \Delta M^2$ gives

$$\mathrm{E}(M_t^2 - \lambda t \mid \mathcal{F}_s) = M_s^2 + 2M_s\,\mathrm{E}(\Delta M) + \mathrm{E}(\Delta M^2) - \lambda t = M_s^2 + \lambda(t-s) - \lambda t = M_s^2 - \lambda s.$$

For the third, the MGF of a Poisson random variable with mean $\lambda(t-s)$ is $\mathrm{E}(e^{\theta\Delta N}) = e^{\lambda(t-s)(e^\theta - 1)}$, so that

$$\mathrm{E}(Z_t \mid \mathcal{F}_s) = Z_s\,\mathrm{E}\left(e^{\theta\Delta N}\right)e^{-\lambda(t-s)(e^\theta-1)} = Z_s.$$

□

These three martingales are the continuous-time counterparts of the examples in the notes on discrete-time martingales. The compensated process plays the role of the random walk S_n, the second martingale plays the role of $S_n^2 - n$, and the exponential martingale has the form $e^{\theta N_t - tK(\theta)}$, where $K(\theta) = \lambda(e^\theta - 1)$ is the cumulant generating function of N_1.

Compound Poisson Processes

In most applications each event carries a size: an insurance claim has an amount, a default causes a loss, and a price jump has a magnitude. Let $Y_1, Y_2, \ldots$ be iid random variables, independent of the Poisson process $\{N_t\}$, and define

$$X_t = \sum_{i=1}^{N_t} Y_i,$$

with $X_t = 0$ when $N_t = 0$. The process $\{X_t\}$ is called a *compound Poisson process*. Its paths are step functions, like those of N_t, but the jumps have random sizes Y_i, which can be positive or negative. A compound Poisson process also has independent and stationary increments.

For each t, X_t is a random sum. The argument used in the notes on generating functions for the PGF of a random sum applies to the characteristic function without change:

given $N_t = n$, the characteristic function of X_t is $\phi_Y(u)^n$, so that $\phi_{X_t}(u) = \mathrm{E}(\phi_Y(u)^{N_t}) = G_{N_t}(\phi_Y(u))$.

Characteristic Function of a Compound Poisson Process

If $\{X_t\}$ is a compound Poisson process with rate λ and jump sizes with characteristic function ϕ_Y, then

$$\phi_{X_t}(u) = e^{\lambda t(\phi_Y(u)-1)}. \tag{1}$$

If $\mathrm{E}(Y^2) < \infty$, then $\mathrm{E}(X_t) = \lambda t\ \mathrm{E}(Y)$ and $\mathrm{V}(X_t) = \lambda t\ \mathrm{E}(Y^2)$. When the MGF of Y exists, $M_{X_t}(\theta) = e^{\lambda t(M_Y(\theta)-1)}$.

Proof

The characteristic function follows from $G_{N_t}(z) = e^{\lambda t(z-1)}$ evaluated at $z = \phi_Y(u)$. For the moments, write $\kappa(u) = \lambda t(\phi_Y(u) - 1)$, so that $\phi_{X_t} = e^{\kappa}$. Differentiating, $\phi'_{X_t} = \kappa' e^{\kappa}$ and $\phi''_{X_t} = (\kappa'' + \kappa'^2)e^{\kappa}$. At $u = 0$ we have $\kappa(0) = 0$, $\kappa'(0) = \lambda t\, i\, \mathrm{E}(Y)$, and $\kappa''(0) = -\lambda t\ \mathrm{E}(Y^2)$. Therefore

$$\mathrm{E}(X_t) = -i\phi'_{X_t}(0) = \lambda t\ \mathrm{E}(Y),$$
$$\mathrm{E}(X_t^2) = -\phi''_{X_t}(0) = \lambda t\ \mathrm{E}(Y^2) + (\lambda t\ \mathrm{E}(Y))^2,$$

and the variance is $\mathrm{E}(X_t^2) - \mathrm{E}(X_t)^2 = \lambda t\ \mathrm{E}(Y^2)$. □

The variance depends on $\mathrm{E}(Y^2)$, not on $\mathrm{V}(Y)$, because the randomness of X_t comes from two sources: the number of jumps and their sizes. Equation (1) is the reason characteristic functions are the natural tool for jump processes. The density of X_t is an infinite mixture over the number of jumps and is rarely available in closed form, whereas its characteristic function is a single exponential. The best-known financial example is the jump-diffusion model of Merton (1976), in which the log price of a stock is the sum of a Brownian component and a compound Poisson process with normally distributed jumps.

Example 3. In the Merton model the jumps Y_i of the log price are normally distributed with mean m and variance s^2. Leaving the Brownian part aside, the characteristic function of the jump component over $[0, t]$ is

$$\phi_{X_t}(u) = \exp\!\left(\lambda t\left(e^{ium - \frac{1}{2}u^2 s^2} - 1\right)\right).$$

Since the Brownian and jump components are independent, the characteristic function of the log price is the product of this expression and the normal characteristic function of the Brownian part. We return to this model in the notes on pricing with jumps. □

Application: Default and Credit Spreads

A simple model of credit risk assumes that a firm defaults at the first event of a Poisson process whose intensity λ is the hazard rate of default, as in the section on memorylessness. The default time τ is then exponential with rate λ, and the probability that the firm survives until time T is

$$\mathrm{P}(\tau > T) = e^{-\lambda T}.$$

Consider a zero-coupon bond issued by the firm that pays one dollar at time T if the firm has not defaulted and nothing otherwise. If the risk-free rate is a constant r and bondholders value the bond at its expected payoff discounted at the risk-free rate, the price of the bond is

$$P^{d}(T) = e^{-rT}\,\mathrm{P}(\tau > T) = e^{-(r+\lambda)T}.$$

The bond is priced as if it were a risk-free bond discounted at the higher rate $r + \lambda$. Its yield exceeds the risk-free rate by exactly λ, so in this model the *credit spread* equals the default intensity.

Discounting expected payoffs at the risk-free rate ignores any compensation investors require for bearing default risk. The notes on pricing with jumps show that the formula remains valid once expectations are computed under the risk-neutral measure, where λ becomes a risk-neutral intensity that is typically larger than the historical default rate. Duffie and Singleton (1999) develop this approach in detail, including random intensities and partial recovery.

Example 4. A firm has a default intensity of $\lambda = 0.02$ per year and the risk-free rate is $r = 0.04$. The probability that the firm survives the next $T = 5$ years is $e^{-0.02 \times 5} = 0.9048$. A 5-year zero-coupon bond issued by the firm is worth

$$P^d(5) = e^{-(0.04+0.02)\times 5} = 0.7408,$$

compared with $e^{-0.04\times 5} = 0.8187$ for a risk-free zero-coupon bond with the same maturity. The yield on the defaultable bond is 6%, a credit spread of 200 basis points. □

Application: The Insurance Ruin Problem

The classic model of ruin in insurance is due to Lundberg and Cramér. An insurance company starts with capital $u > 0$ and collects premiums at a constant rate c per unit of time. Claims arrive as a Poisson process with rate λ, and the claim sizes $Y_1, Y_2, \ldots$ are iid, positive, and independent of the arrival times. The surplus of the company at time t is

$$U_t = u + ct - X_t, \qquad X_t = \sum_{i=1}^{N_t} Y_i.$$

The company is ruined at the first time τ its surplus becomes negative, and the ruin probability is $\psi(u) = \mathrm{P}(\tau < \infty)$. On average the surplus grows at the rate $c - \lambda\,\mathrm{E}(Y)$, which we assume is positive. Otherwise, ruin is certain.

As Figure 4 shows, the surplus rises steadily between claims and falls only when a claim arrives. Ruin therefore happens at a claim larger than the remaining surplus, and the surplus lands below zero by the difference. How far below zero it lands will matter for the ruin probability.

Whether a path is ever ruined depends on its whole future, so $\psi(u)$ cannot be read off the distribution of U_t at any single date. As for the gambler in the notes on discrete-time martingales, we use a martingale and compare its value at the start with its value at ruin.

An Exponential Martingale

For the biased gambler in the notes on discrete-time martingales, the martingale was $(q/p)^{S_n}$, a constant raised to the power of the gambler's wealth. We look for a martingale of the same kind, $Z_t = e^{-RU_t}$ for some $R > 0$, with e^{-R} in the role of q/p. Using the MGF of the compound Poisson process,

$$\mathrm{E}(Z_t) = Z_0\, e^{-cRt}\, \mathrm{E}\left(e^{RX_t}\right) = Z_0 \exp(t(\lambda(M_Y(R) - 1) - cR)).$$

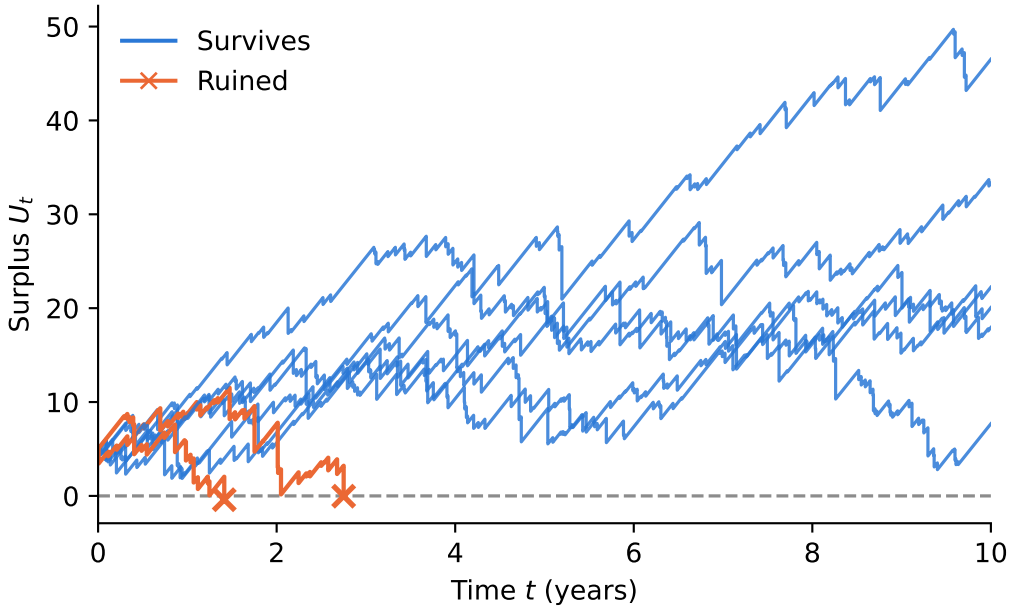


Figure 4: The figure plots eight simulated surplus paths over ten years for an insurer that starts with $u = 5$, receives $\lambda = 10$ claims per year with mean 1, and collects premiums at the rate $c = 12$. Paths that fall below zero are ruined and stop at the ruin time, marked with a cross.

Premiums pull the exponent down at the rate cR, and claims push it up at the rate $\lambda(M_Y(R) - 1)$. The expectation stays at Z_0 when the two cancel,

$$\lambda(M_Y(R) - 1) = cR, \tag{2}$$

and Z is then a martingale, because the increments of X are independent of the past. A positive solution R is called the *adjustment coefficient*. Both sides of (2) are zero at $R = 0$, and the left side is convex with initial slope $\lambda\,\mathrm{E}(Y) < c$, so it starts below the line cR. It crosses the line at some $R > 0$ when $M_Y(\theta)$ grows without bound as θ approaches the end of the interval where it is finite: infinity for bounded claims, and the rate parameter for exponential or gamma claims. Heavy-tailed claims, such as Pareto losses from natural catastrophes, have an infinite MGF for every $\theta > 0$ and no adjustment coefficient.

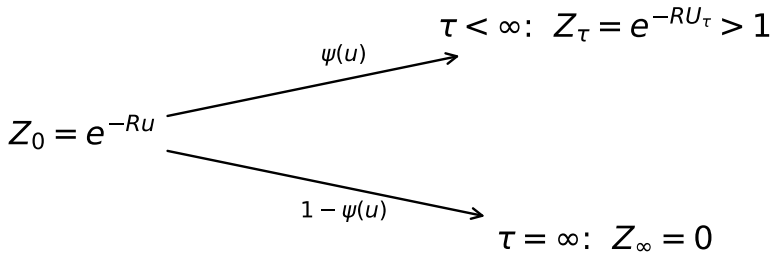


Figure 5: The figure shows the martingale $Z_t = e^{-RU_t}$ stopped at ruin. With probability $\psi(u)$ the company is ruined and Z_τ is larger than one; with probability $1 - \psi(u)$ it is never ruined and Z_t goes to zero.

Stopping at Ruin

The martingale starts at $Z_0 = e^{-Ru}$. On a path that is never ruined, the surplus grows without bound and Z_t goes to zero. On a path that is ruined, the surplus at ruin is negative and $Z_\tau = e^{-RU_\tau}$ is larger than one. The martingale keeps its expectation, so

$$e^{-Ru} = (1 - \psi(u)) \times 0 + \psi(u)\, \mathrm{E}\,(Z_\tau \mid \tau < \infty).$$

Figure 5 draws this calculation as a tree with two branches, like the binomial trees that price options. The value today is the probability-weighted average of the values on the two branches, and the ruin probability is the weight on the branch where the company is ruined.

Ruin Probability

If $R > 0$ solves (2), then for every initial capital $u > 0$,

$$\psi(u) = \frac{e^{-Ru}}{\mathrm{E}\,(e^{-RU_\tau} \mid \tau < \infty)} \leq e^{-Ru}. \tag{3}$$

The bound is called *Lundberg's inequality*.

Proof

The argument above stops the martingale at a time that may be infinite, which needs care: in the notes on discrete-time martingales, a gambler on a fair coin with no upper barrier is ruined with certainty, yet the walk stopped at ruin does not keep its expectation. The proof follows the argument given there for a biased coin. The stopped process $Z_{t\wedge\tau}$ is a martingale, a result that carries over to continuous time, so for any t

$$e^{-Ru} = \mathrm{E}(Z_{t\wedge\tau}) = \mathrm{E}(Z_\tau \mathbf{1}_{\{\tau\le t\}}) + \mathrm{E}(Z_t \mathbf{1}_{\{\tau>t\}}).$$

Let $t \to \infty$. On each path, $Z_\tau \mathbf{1}_{\{\tau\le t\}}$ is zero before ruin and Z_τ from ruin on, so it increases with t, and by monotone convergence the first term goes to $\mathrm{E}(Z_\tau \mathbf{1}_{\{\tau<\infty\}})$. For the second term,

$$\frac{U_t}{t} = \frac{u}{t} + c - \frac{N_t}{t} \cdot \frac{1}{N_t} \sum_{i=1}^{N_t} Y_i \to c - \lambda\,\mathrm{E}(Y) > 0$$

almost surely, by the strong law of large numbers applied to the claims and to the interarrival times, whose average $1/\lambda$ gives $N_t/t \to \lambda$. Hence $U_t \to \infty$ and $Z_t \mathbf{1}_{\{\tau>t\}} \to 0$ almost surely. Before ruin the surplus is nonnegative, so $Z_t \mathbf{1}_{\{\tau>t\}} \le 1$, and by dominated convergence $\mathrm{E}(Z_t \mathbf{1}_{\{\tau>t\}}) \to 0$. Therefore $e^{-Ru} = \mathrm{E}(Z_\tau \mathbf{1}_{\{\tau<\infty\}}) = \psi(u)\,\mathrm{E}(Z_\tau \mid \tau < \infty)$. □

The denominator in (3) depends on how far below zero the surplus falls at ruin. The gambler moves in unit steps and lands exactly at zero, so for the gambler the denominator is one and the formula is exact: for a walk that drifts up, with $p > q$ and no upper barrier, it gives a ruin probability of $(q/p)^u$. The insurer's surplus jumps past zero by an amount that depends on the claims, but in every case Lundberg's inequality shows that the ruin probability falls exponentially with the capital.

The Deficit with Exponential Claims

When claims are exponential with mean $1/\eta$, the deficit at ruin can be computed. Let V be the surplus just before the claim that causes ruin, so that $U_\tau = V - Y_k$ when ruin happens at the k-th claim. Whether the company survives the first $k-1$ claims, and the surplus v it has just before the k-th, depend only on earlier claims and arrival times, so they are independent of Y_k. On that event, ruin at the k-th claim means $Y_k > v$. For $y \ge 0$,

$$\mathrm{P}(U_\tau < -y \mid \text{ruin at claim } k, V = v) = \mathrm{P}(Y_k > v + y \mid Y_k > v) = \mathrm{P}(Y_k > y) = e^{-\eta y},$$

where the second step is the memoryless property of the exponential distribution. The answer depends neither on k nor on v, so averaging over both by the law of total probability gives

$$\mathrm{P}(U_\tau < -y \mid \tau < \infty) = e^{-\eta y}.$$

Given ruin, the deficit $-U_\tau$ has the same distribution as a claim, however close to zero the company was.

Example 5. With exponential claims, $M_Y(\theta) = \eta/(\eta - \theta)$ for $\theta < \eta$, and (2) becomes $\lambda R/(\eta - R) = cR$, with positive solution $R = \eta - \lambda/c$. Given ruin, $U_\tau \sim -Y$, and

$$\mathrm{E}\left(e^{-RU_\tau} \mid \tau < \infty\right) = \mathrm{E}\left(e^{RY}\right) = \frac{\eta}{\eta - R} = \frac{c\eta}{\lambda}.$$

Substituting into (3),

$$\psi(u) = \frac{\lambda}{c\eta} e^{-(\eta - \lambda/c)u}.$$

The ruin probability falls with the capital u, which enters only through the exponential decay. It also falls with the premium rate c and rises with the claim rate λ and the mean claim $1/\eta$, each of which moves the factor and the decay rate in the same direction.

The insurer in Figure 4 receives $\lambda = 10$ claims per year with a mean size of 1 million dollars and collects premiums of $c = 12$ million dollars per year, 20% above expected claims. The adjustment coefficient is $R = 1 - 10/12 = 1/6$. With an initial capital of $u = 5$ million dollars, Lundberg's bound is $e^{-5/6} = 0.4346$ and the exact ruin probability is $\frac{10}{12} e^{-5/6} = 0.3622$. A simulation of 4000 surplus paths over about 150 years, long enough that later ruin is negligible, yields a ruin frequency of 0.364. To bring the ruin probability down to 1%, the insurer needs capital of $u = 6 \ln(100 \times 10/12) = 26.5$ million dollars. □

Practice Problems

Problem 1. Orders reach a market maker as a Poisson process with rate $\lambda = 3$ per minute.

a. Compute the probability that no order arrives during a given minute, and the probability that at least two arrive.
b. Compute the mean and standard deviation of the waiting time until the fifth order.
c. Exactly one order arrived during the first minute. What is the probability that exactly four orders have arrived by the end of the second minute?

Problem 2. A firm defaults at the first event of a Poisson process with intensity $\lambda = 0.03$ per year.

a. Compute the probability that the firm defaults within five years.
b. The firm has survived for ten years. Compute the probability that it defaults within the next five years.
c. Compute the median and the mean of the default time, and the probability that the firm survives beyond the mean. Why is the median below the mean?

Problem 3 (⋆). Three firms default at independent exponential times τ_1, τ_2, τ_3 with intensities $\lambda_1 = 0.01$, $\lambda_2 = 0.02$, and $\lambda_3 = 0.03$ per year.

a. Show that the first default time $\min(\tau_1, \tau_2, \tau_3)$ is exponential with rate $\lambda_1 + \lambda_2 + \lambda_3$, and compute its mean.
b. Show that $\mathrm{P}(\tau_3 < \min(\tau_1, \tau_2)) = \lambda_3/(\lambda_1 + \lambda_2 + \lambda_3)$. Hint: condition on τ_3.
c. A *first-to-default* bond pays one dollar in five years if none of the three firms has defaulted. With a risk-free rate of 4%, price it as in the notes, and compare it with a five-year bond of the riskiest firm alone.

Problem 4 (⋆). Let $\{N_t\}$ be a Poisson process with rate λ, and let $0 < s < t$.

a. Show that, given $N_t = n$, N_s has the binomial distribution with n trials and success probability s/t. Hint: write $N_t = N_s + (N_t - N_s)$.
b. Conclude that, given $N_t = 1$, the time of the event is uniformly distributed on $[0, t]$.
c. A loan book had four defaults last year. Compute the probability that all four occurred in the first half of the year, and the expected number in the first quarter. Does the answer depend on λ?

Problem 5 (★). Leave the Brownian part of the Merton model aside, so the log price over one year is $X_1 = \sum_{i=1}^{N_1} Y_i$, with $\lambda = 2$ jumps per year and jumps $Y_i \sim \mathcal{N}(m, s^2)$, where $m = -0.05$ and $s = 0.10$.

a. Compute $\mathrm{E}(X_1)$, $\mathrm{SD}(X_1)$, and the probability that the log price does not move during the year.
b. Rederive $\mathrm{V}(X_t) = \lambda t\, \mathrm{E}(Y^2)$ from $\mathrm{V}(X_t) = \mathrm{E}(\mathrm{V}(X_t \mid N_t)) + \mathrm{V}(\mathrm{E}(X_t \mid N_t))$.
c. Show that the cumulant generating function of X_t is $\lambda t(M_Y(\theta) - 1)$, so that its k-th cumulant is $\lambda t\, \mathrm{E}(Y^k)$. Compute the skewness of X_1.

Problem 6 (★). Let W_n be the time of the n-th event of a Poisson process with rate λ, and let $r > 0$ be the interest rate.

a. Use the MGF of W_n to show that the present value of one dollar paid at time W_n is $\mathrm{E}(e^{-rW_n}) = (\lambda/(\lambda + r))^n$.
b. A reinsurer must pay \$1 million at the third major hurricane. Hurricanes arrive at rate $\lambda = 0.1$ per year and $r = 0.04$. Compute the present value of the obligation.
c. A perpetual bond pays a coupon at the rate of \$5 per year until the issuer defaults at the first event of a Poisson process with intensity $\lambda = 0.02$, and nothing afterwards. Show that its value is $c/(r + \lambda)$, where c is the coupon rate, and compute it for $r = 0.04$. Compare with the risk-free perpetuity.

Problem 7 (★★). A stock price moves by one tick at the events of a Poisson process with rate λ, up or down with probability $1/2$ each, independently of the event times. Its change since time zero is the compound Poisson process $X_t = \sum_{i=1}^{N_t} Y_i$, with $Y_i = \pm 1$.

a. Use thinning to write $X_t = N_t^u - N_t^d$ as the difference of two independent Poisson processes, and compute $\mathrm{E}(X_t)$ and $\mathrm{V}(X_t)$.
b. Show that X_t and $X_t^2 - \lambda t$ are martingales, and that $\exp(\theta X_t - \lambda t(\cosh\theta - 1))$ is a martingale for every θ.
c. Explain how X_t relates to the simple symmetric random walk, and compare the three martingales with those in the notes on discrete-time martingales.

Problem 8 (★★). An insurer receives claims at rate $\lambda = 1$ per year and collects premiums at the rate $c = 1.2$ per year, in millions of dollars.

a. Every claim is exactly \$1 million. Write the equation for the adjustment coefficient R, solve it numerically, and compute Lundberg's bound on the ruin probability for an initial capital of $u = 10$.
b. Claims are instead exponential with mean \$1 million. Compute R, the bound, and the exact ruin probability for $u = 10$.
c. Use $e^x \geq 1 + x + \frac{1}{2}x^2$ for $x \geq 0$ to show that for any positive claims

$$R \leq \frac{2(c - \lambda\,\mathrm{E}(Y))}{\lambda\,\mathrm{E}(Y^2)}.$$

Evaluate the bound in both cases, and explain why two claim distributions with the same mean lead to different capital requirements.

References

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

Grimmett, Geoffrey, and David Stirzaker. 2001. *Probability and Random Processes*. 3rd ed. Oxford University Press.

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.