

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 plays the same role for sudden events that Brownian motion plays for continuous noise. Both arise as limits of sums of Bernoulli random variables, both have independent and stationary increments, and both generate martingales that make calculations tractable. The difference lies in the paths. Brownian motion moves all the time by tiny amounts, whereas a Poisson process stays constant most of the time and moves only by jumps. Studying the Poisson process first, before we build Brownian motion, lets us see these ideas in a setting where the paths are simple to draw and the calculations require nothing beyond generating functions and martingales.

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

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

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 $0 \leq t_0 < t_1 < \dots < t_n$, the increments $N_{t_1} - N_{t_0}, N_{t_2} - N_{t_1}, \dots, N_{t_n} - N_{t_{n-1}}$ are independent.
- iii. For any $0 \leq s < t$, the increment $N_t - N_s$ has a Poisson distribution with mean $\lambda(t - s)$.

Compare this definition with the definition of Brownian motion. The first two properties are identical, and the third replaces a normal increment with variance $t - s$ by a Poisson increment 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

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

The rate λ is the expected number of events per unit of time. A firm with a default intensity of $\lambda = 0.02$ per year, for example, expects on average 0.02 defaults per year, or one default every 50 years.

The Bernoulli construction also describes the behavior of the process over a short interval. Over an interval of length dt , the Poisson probabilities give

$$P(N_{t+dt} - N_t = 0) = e^{-\lambda dt} = 1 - \lambda dt + o(dt),$$

$$P(N_{t+dt} - N_t = 1) = \lambda dt e^{-\lambda dt} = \lambda dt + o(dt),$$

$$P(N_{t+dt} - N_t \geq 2) = o(dt),$$

where $o(dt)$ denotes terms that vanish faster than dt . 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 $dN_t = N_{t+dt} - N_t$, this says that dN_t is approximately a Bernoulli random variable with $E(dN_t) = \lambda dt$.




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 $E(N_t) = \lambda t$.

Waiting Times

Interarrival Times Are Exponential

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

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

Therefore, T_1 has an exponential distribution with rate λ , and the expected waiting time for the first event is $1/\lambda$.

The Bernoulli construction gives the same answer from a different angle. In discrete time, the number of slots until the first event is geometric with success probability $\lambda\Delta t$, so the probability that no event occurs in the first $t/\Delta t$ slots is

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

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.

The same reasoning applies after the first event. Because increments are independent and stationary, once an event has occurred the process starts afresh, 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 $T_1, T_2, \dots$ denote the times between consecutive events of a Poisson process with rate λ . Then $T_1, T_2, \dots$ are independent random variables, each with an $\text{Exp}(\lambda)$ distribution. Conversely, if $T_1, T_2, \dots$ are iid $\text{Exp}(\lambda)$ and $W_n = T_1 + \dots + T_n$, 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.

Memorylessness

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

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

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 $\{1, 2, \dots\}$.

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

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

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.

Arrival Times

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

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

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

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

Example 1. 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, using the superposition property discussed below. Ignoring the small reduction in the rate after each default, the probability that at least two firms default within a year is

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

□

Martingales of the Poisson Process

Let $\mathcal{F}_t$ denote the information generated by observing the process $\{N_s : s \leq t\}$ 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 $E(M_t | \mathcal{F}_s) = M_s$ for all $s \leq t$. 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 $Z_t = \exp(\theta N_t - \lambda t(e^\theta - 1))$.

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 $E(\Delta M) = 0$ and $E(\Delta M^2) = V(\Delta N) = \lambda(t - s)$.

For the first process,

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

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

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

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

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

□

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 . We will soon see that Brownian motion has three martingales with exactly the same structure.

Quadratic Variation

For a partition $0 = t_0 < t_1 < \dots < t_n = t$, consider the sum of squared increments of the Poisson process,

$$Q_\Pi = \sum_{j=0}^{n-1} (N_{t_{j+1}} - N_{t_j})^2.$$

There are finitely many events in $[0, t]$, and no two occur at the same time. As the mesh of the partition goes to zero, each subinterval eventually contains at most one event, so each increment is either zero or one, and its square equals itself. Therefore, the sum of squared increments equals the sum of increments, and

$$[N, N]_t = \lim_{\|\Pi\| \rightarrow 0} Q_\Pi = N_t.$$

The quadratic variation of a Poisson process counts its jumps. The same argument applies to the compensated process $M_t = N_t - \lambda t$: subtracting the smooth function λt does not change the quadratic variation, so $[M, M]_t = N_t$ as well.

In differential notation this result reads $(dN)^2 = dN$, which together with $(dN)(dt) = 0$ and $(dt)^2 = 0$ gives the multiplication table for the Poisson process. Compare it with the result we will prove for Brownian motion, $[B, B]_t = t$, or $(dB)^2 = dt$:

- The quadratic variation of Brownian motion is deterministic. The quadratic variation of the compensated Poisson process is random, and only its expectation, $E(N_t) = \lambda t$, is deterministic.
- In both cases the expected quadratic variation equals the variance of the process, $V(B_t) = t$ and $V(M_t) = \lambda t$.
- Brownian paths accumulate quadratic variation continuously through infinitely many tiny moves. Poisson paths accumulate it only at the jumps.

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

- 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 $\lambda_1 + \lambda_2$.
- 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 $\lambda(1 - p)$.

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 $(\lambda_1 + \lambda_2)(t - s)$.

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

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

Taking expectations over N_t and using its PGF,

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

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

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.

Example 2. 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. $\square$

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 $Y_1, Y_2, \dots$ 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 $\phi_{X_t}(u) = E(\phi_Y(u)^{N_t}) = G_{N_t}(\phi_Y(u))$.

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

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

If $E(Y^2) < \infty$, then $E(X_t) = \lambda t E(Y)$ and $V(X_t) = \lambda t E(Y^2)$. 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 $\psi(u) = \lambda t(\phi_Y(u) - 1)$, so that $\phi_{X_t} = e^\psi$. Differentiating, $\phi'_{X_t} = \psi' e^\psi$ and $\phi''_{X_t} = (\psi'' + \psi'^2) e^\psi$. At $u = 0$ we have $\psi(0) = 0$, $\psi'(0) = \lambda t i E(Y)$, and $\psi''(0) = -\lambda t E(Y^2)$. Therefore

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

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

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

The variance depends on $E(Y^2)$, not on $V(Y)$, 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

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

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. $\square$

Application: Default and Credit Spreads

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

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

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

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

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 five-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. $\square$

Application: The Insurance Ruin Problem

The ruin problem for a gambler has a classic continuous-time counterpart in insurance, the model of Lundberg and Cramér. An insurance company starts with capital $u > 0$ and collects premiums continuously at a rate of c per unit of time. Claims arrive as a Poisson process with rate λ , and the claim sizes $Y_1, Y_2, \dots$ are iid, positive, and independent of the arrival times. The surplus of the company at time t is

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

The company is ruined at the first time τ its surplus becomes negative, and we want the ruin probability $\psi(u) = P(\tau < \infty)$. The surplus increases at a constant rate $c - \lambda E(Y)$ on average, and we assume this rate is positive, $c > \lambda E(Y)$. Otherwise, ruin is certain.

Unlike the gambler, the company faces no upper barrier, and the surplus falls only at claim times, by random amounts. Exact ruin probabilities are hard to compute in general, but an exponential martingale gives a sharp bound.

Lundberg's Inequality

Suppose there is a number $R > 0$, called the *adjustment coefficient*, that solves

$$\lambda(M_Y(R) - 1) = cR. \quad (2)$$

Then the ruin probability satisfies $\psi(u) \leq e^{-Ru}$ for every initial capital $u > 0$.

Proof

First we show that $Z_t = e^{-Ru_t}$ is a martingale. The increment $U_t - U_s = c(t - s) - (X_t - X_s)$ is independent of $\mathcal{F}_s$, and $X_t - X_s$ has the same distribution as X_{t-s} . Using the MGF of the compound Poisson process,

$$\mathbb{E}(Z_t | \mathcal{F}_s) = Z_s e^{-Rc(t-s)} \mathbb{E}(e^{R(X_t - X_s)}) = Z_s \exp((t - s)(\lambda(M_Y(R) - 1) - cR)) = Z_s,$$

where the last step uses (2).

As in the notes on discrete-time martingales, the stopped process $Z_{t \wedge \tau}$ is also a martingale, a result that carries over to continuous time. Therefore, for any t ,

$$e^{-Ru} = Z_0 = \mathbb{E}(Z_{t \wedge \tau}) \geq \mathbb{E}(Z_\tau \mathbf{1}_{\{\tau \leq t\}}),$$

where the inequality holds because Z is positive. At the ruin time the surplus is negative, so $Z_\tau = e^{-RU_\tau} > 1$, and hence $e^{-Ru} \geq \mathbb{P}(\tau \leq t)$. Letting $t \rightarrow \infty$ gives $e^{-Ru} \geq \mathbb{P}(\tau < \infty) = \psi(u)$. $\square$

The structure of the proof is the same as for the gambler: find a martingale, stop it at the ruin time, and compare the value at the start with the value at the stopping time. Here the martingale is exponential in the surplus. For a biased random walk that drifts up, with $p > q$ and no upper barrier, the same argument applied to the martingale $(q/p)^{S_n}$ gives a ruin probability of $(q/p)^u$ from an initial wealth of u , and e^{-R} plays the role of q/p . The bound shows that the ruin probability falls exponentially fast with the initial capital.

Example 5. Suppose claims are exponentially distributed with mean $1/\eta$, so that $M_Y(\theta) = \eta/(\eta - \theta)$ for $\theta < \eta$. Equation (2) becomes

$$\lambda \left(\frac{\eta}{\eta - R} - 1 \right) = \frac{\lambda R}{\eta - R} = cR,$$

with positive solution $R = \eta - \lambda/c$. An insurer 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 inequality gives

$$\psi(5) \leq e^{-5/6} = 0.4346.$$

For exponential claims the exact ruin probability is known to be $\psi(u) = \frac{\lambda}{c\eta} e^{-Ru}$, which equals 0.3622 here. A simulation of 4000 surplus paths over about 150 years yields a ruin frequency of 0.364, close to the exact value. $\square$




Figure 2: The figure plots eight simulated surplus paths for the insurer in [Example 5](#) over ten years, starting from $u = 5$. Paths that fall below zero are ruined and stop at the ruin time, marked with a cross.

Summary

The table below compares the Poisson process with Brownian motion, which we construct in the next notes.

Property	Poisson process N_t	Brownian motion B_t
Increment over $(s, t]$	Poisson, mean $\lambda(t - s)$	Normal, variance $t - s$
Discrete-time origin	Rare Bernoulli events	Scaled coin flips
Paths	Step functions, jumps of size one	Continuous, nowhere differentiable
Martingale	$N_t - \lambda t$	B_t
Second martingale	$(N_t - \lambda t)^2 - \lambda t$	$B_t^2 - t$
Exponential martingale	$e^{\theta N_t - \lambda t(e^\theta - 1)}$	$e^{\theta B_t - \frac{1}{2}\theta^2 t}$
Quadratic variation	N_t (random)	t (deterministic)

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.