Birth and Poisson processes

Olivier Scaillet

University of Geneva and Swiss Finance Institute

Outline

Introduction

Many natural phenomenons exhibit changes of value at random times rather than fixed times.

We model this behavior with continuous-time models, that is, processes of the form $\{X_t : t \geq 0\}$, where t is continuous.

The sample paths of the processes we consider here have integer values.

Poisson process

The Poisson process is ideal to model the occurrence of successive events.

Every event is such that, in a time interval $(t, t + \Delta t)$ with Δt small, its probability of occurrence is proportional to Δt, and that the probability of occurrence of two events in the same small interval is negligible.

Poisson process (cont'd)

A Poisson process with intensity λ is a process $N = \{N_t : t \geq 0\}$, whose values belong to $0, 1, 2, 3, \ldots$ and such that:

1. $N_0 = 0$ and $N_s \leq N_t$, if $s < t$,
2. if $s < t$, the number $N_t - N_s$ of events in the time interval $(s, t]$ is independent of the arrival times and the arrivals in the interval $[0, s]$ (process with independent increments)
3. $P[N_{t+\Delta t} = n + i | N_t = n] = \begin{cases} \lambda \Delta t + o(\Delta t) & \textit{if } i = 1 \\ o(\Delta t) & \textit{if } i > 1 \\ 1 - \lambda \Delta t + o(\Delta t) & \textit{if } i = 0 \end{cases}$

Notation: $f(\Delta t) = o(\Delta t)$ is equivalent to $\lim_{\Delta t \to 0} \frac{f(\Delta t)}{\Delta t} = 0$.

Poisson process (cont'd)

The function f has a smaller asymptotic order than Δt, and is negligible in comparison to Δt.

We say that N_t stands for the number of arrivals, occurrences, or events of the process until time t. The process $\{N_t\}$ is a counting process and is a simple example of continuous-time Markov chains.

The intensity λ (or birth rate) corresponds to the probability of occurrence of one event per unit of time.

The Poisson process N_t follows a Poisson distribution with parameter λt, $P[N_t = i] = \frac{(\lambda t)^i}{i!} e^{-\lambda t}$, $i = 0, 1, 2, \ldots$. In particular, we have $E[N_t] = \lambda t$.

Poisson process (cont'd)

We can also characterize the Poisson process using the duration between successive events.

Let $T_0, T_1, \dots$ be given by $T_0 = 0$, $T_n = \inf\{t : N_t = n\}$. T_n is the date of the nth arrival.

The durations (or intervals) are random variables $d_n = T_n - T_{n-1}$. These variables $d_1, d_2, \dots$ are independent and follow an exponential distribution with parameter λ, whose density function is given by $f(d) = \lambda e^{-\lambda d}$.

The distribution function is given by $F(d) = 1 - e^{-\lambda d}$, and the expectation of the durations is $E[d_t] = \frac{1}{\lambda}$.

Poisson process (cont'd)

Remarks:

- one can rebuild the process N with durations d, using $T_n = \sum_{i=1}^{n} d_i$, $N_t = \max\{n : T_n \leq t\}$.
- the estimation of the lambda parameter is trivial as the maximum likelihood estimator for an exponential law is the inverse of the empirical mean:$\hat{\lambda} = \frac{1}{\bar{d}}$, where $\bar{d}$ is the empirical mean of the observed durations.

Birth process

The Poisson process has a constant intensity (or birth rate). We may introduce a dependence between the number of previous events and the birth rate.

A birth process with intensities $\lambda_0, \lambda_1, \lambda_2, \dots$ is a process $N = \{N_t : t \geq 0\}$, whose values lie in $\{0, 1, 2, 3, \dots\}$ and such that:

1. $N_0 = 0$ and $N_s \leq N_t$, if $s < t$
2. if $s < t$, conditionally on the value of N_s, the number of events $N_t - N_s$ in the time interval $(s, t]$ is independent of all previous arrivals in the interval$[0, s]$
3. $P[N_{t+\Delta t} = n + i | N_t = n] = \begin{cases} \lambda_n \Delta t + o(\Delta t) & \textit{if } i = 1 \\ o(\Delta t) & \textit{if } i > 1 \\ 1 - \lambda_n \Delta t + o(\Delta t) & \textit{if } i = 0 \end{cases}$

Birth process (cont'd)

Particular cases:

- Poisson process: $\lambda_n = \lambda$, $\forall n$
- Simple birth: $\lambda_n = n\lambda$. Population growth where each individual gives birth independently of the others, with probability $\lambda \Delta t + o(\Delta t)$ during the time interval $(t, t + \Delta t)$, and where there is no death.
- Birth with immigration: $\lambda_n = n\lambda + v$. Birth process where the birth rate depends on an external immigration rate.

Birth and death process

A more realistic model in the context of a population growth also allows for deaths.

Let N_t be the number of alive individuals at time t in a particular population. This number follows a birth and death process if:

1. N_t is a Markov chain with values $\{0, 1, 2, \ldots\}$
2. $P[N_{t+\Delta t} = n + i | N_t = n] =$
$$\begin{cases} \lambda_n \Delta t + o(\Delta t) & \text{if } i = 1 \\ \mu_n \Delta t + o(\Delta t) & \text{if } i = -1 \\ o(\Delta t) & \text{if } |i| > 1 \\ 1 - \lambda_n \Delta t - \mu_n \Delta t + o(\Delta t) & \text{if } i = 0 \end{cases}$$
3. The birth rates $\lambda_0, \lambda_1, \lambda_2, \ldots$ and death rates $\mu_0, \mu_1, \mu_2, \ldots$ are such that $\lambda_i \geq 0$, $\mu_i \geq 0$, and $\mu_0 = 0$.

Birth and death process (cont'd)

Particular cases:

- Birth process: $\mu_n = 0$, $\forall n$.
- Simple death with immigration: $\lambda_n = \lambda$, $\mu_n = n\mu$, $\forall n$. Population that does not give birth but with some immigration corresponding to a Poisson process, and for which each individual has a probability $\mu \Delta t + o(\Delta t)$ to die in the time interval $(t, t + \Delta t)$.
- Simple birth and death: $\lambda_n = n\lambda$, $\mu_n = n\mu$, $\forall n$. Each invididual that is alive may either die during the time interval $(t, t + \Delta t)$ with a probability $\mu \Delta t + o(\Delta t)$, or split with a probability $\lambda \Delta t + o(\Delta t)$.

Waiting queues

Waiting queues theory is an important application of the stochastic process theory.

Practical examples:

- Connections to a server.
- Call center capacity.
- Number of counters in a bank.

Inflows

A waiting queue has an input stream that stands for the client arrivals ("client" has a broad meaning, e.g., cars, memory calls of a computer, phone call in a call center).

As a first approximation, we assume that the durations between arrivals are i.i.d. random variables.

The input stream is then a so-called stationary renewal process. We get the simplest case by using a duration following the exponential law: the input flow follows a Poisson process.

Service

A waiting queue generally has a service provider that is specified by

- The duration of the service, whose distribution is assumed to be known.
- The number of ticket offices (or counters).

Service (cont'd)

A waiting queue is also characterized by a rule of service that defines the service behavior:

- System with or without waiting queue (i.e., in a system without queue, clients are lost if they are not served immediately); system with a limited capacity (i.e., clients are lost when the queue is too long).
- The order in which clients are served: first in first out (FIFO), last in first out (LIFO).
- Multiple client priority classes.
- Contingent probability of client loss: the client may be lost with a probability depending on the length of the queue.

Kendall notation

For a mathematical study, we classify the waiting queues following the standardized Kendall notation A/B/C/D. The letters have the following meaning:

- A: input stream
- B: duration of service
- C: number of ticket offices
- D: additional information, as follows.

The letter M (for Markov) for input streams stands for a Poisson stream; for the duration of service, it stands for an exponential duration.

The letter D (for Determinist) for input streams stands for regular interval durations; for the duration of service, it stands for a fixed time for any client.

The letter G stands for the general case.

Example: M/M/1 queue

It is the simplest example of a waiting queue:

- The input stream follows a Poisson process with intensity λ, or, equivalently, the durations between the arrivals follow an i.i.d. exponential distribution with parameter λ.
- The service durations are i.i.d. and follow an exponential law with parameter μ.
- There is only one ticket office and client are served according to a first in first out rule; the capacity of the queue is infinite.

Remark: the M/M/1 queue is equivalent to a birth and death process whose birth and death rates are given by $\lambda_n = \lambda$, $\forall n$ and $\mu_0 = 0$, $\mu_n = \mu$, $\forall n > 0$, respectively.

Example: M/M/1 queue (cont'd)

If the ticket office is open without interruption, the average service rate is equal to μ. As the average input flow is λ per unit of time, it is unsustainable if $\lambda > \mu$.

We measure the traffic intensity with the ratio $\rho = \frac{\lambda}{\mu}$ and use it to analyse the behavior of the queue.

Let N_t be the number of clients in the system at time t (i.e., clients in the queue or being served):

- The probability that a client being served at time t finishes at time $t + \Delta t$ does not depend on t (exponential service).
- The probability that a new client arrives during the time interval $(t, t + \Delta t)$ does not depend on t (exponential arrivals).

Example: M/M/1 queue (cont'd)

We can show that if $\rho < 1$ (or equivalently $\lambda < \mu$), that is, the average input flow is strictly smaller than the service flow.

- The random number of people in the process N_t is a recurrent process: it can take values that are arbitrarily large or the value zero an infinite number of times.
- There is a stationary regime given by the the law $P[N_t = n] \to \pi_n = (1-\rho)\rho^n$ as $t \to \infty$.

We deduce that the average number of people in the system is $E^{\pi}[N_t] = \sum_n n\pi_n = \frac{\rho}{1-\rho}$.

Example: M/M/1 queue (cont'd)

If $\rho > 1$ (or eq. $\lambda > \mu$), that is, the average input flow is larger than the service flow, the length of the queue goes to infinity.

If $\rho = 1$ (or eq. $\lambda = \mu$), that is, if the average input flow is equal to the service flow, the random number of people in the system N_t is a recurrent system but there is no stationary regime (e.g., there are large oscillations in the number of people in the waiting queue).

Exemple: M/M/1 queue (cont'd)

Remarks:

- We can generally characterize the laws and moments of the queue length, that is, the number of people waiting to be served, the closing times required (while maintaining the service) to reduce the queue, etc.
- We may also analyse the queues M/M/k, M/M/∞, ..., or impose bounds on the arrival and waiting times, work with simultaneous arrivals (batches), etc.

Risk process

In insurance, we use risk processes in order to determine the amount of reserves required to finance occasional sinisters (or claims). These risk models use Poisson processes to model the arrival of sinisters.

Let $N = \{N_t\}$ be a Poisson process with constant intensity λ and let $U_1, U_2, \ldots$ be i.i.d. variables.

The process $N^* = \{N_t^* = U_1 + U_2 + \cdots + U_{N_t}\}$ is called compound Poisson process. U_n is the change of N_t^*.

Risk process (cont'd)

At the time of the n^{th} arrival of the Poisson process, a random quantity U_n is added to N_t^*.

The random variables $U_1, U_2, \ldots$ stand for the sinisters and the Poisson process models their arrival times. The durations between services thus follow an exponential law with parameter λ.

Reserve risk model

The reserve risk model based on a compound Poisson process is made of:

1. Random times at which sinisters are recorded, and described by a Poisson process $\{N_t\}$ with constant intensity λ.
2. Rhe sizes of the respective sinisters $U_1, U_2, \ldots$ are i.i.d. and positive by convention.
3. An initial risk reserve u.
4. The premium are assumed to be gathered at a constant rate $\beta > 0$ such that the return is a linear function of time.

The reserve risk process $\{R_t : t \geq 0\}$ is then given by $R_t = u + \beta t - \sum_{i=1}^{N_t} U_i$.

The surplus process $\{S_t : t \geq 0\}$ is given by $S_t = \sum_{i=1}^{N_t} U_i - \beta t$.

Ruin

The time of ruin $\tau(u) = \min\{t : R_t < 0\} = \min\{t : S_t > u\}$ is the first time at which the reserve process gets negative, or equivalently, the surplus process gets larger than u.

We generally study the general behavior of the ruin probability

1. at a finite horizon $\psi(u; x) = P[\tau(u) \leq x]$
2. at an infinite horizon $\psi(u) = \lim_{x\to\infty} \psi(u; x) = P[\tau(u) \leq \infty]$.

The probability $\psi(u)$ corresponds to an ultimate probability of ruin, and $\bar{\psi} = 1 - \psi(u)$ to a survival probability.

Cat bond

Cat bonds are used as alternative financing sources in insurance.

It is based on the securitization of the catastrophe risk.

Why would we need alternative sources in addition to the usual reinsurance process?

Cat bond (cont'd)

1. Population growth in regions with high risk of disasters such as earthquakes or tornadoes (e.g., growth in California during the last 30 years is three times larger than the average growth in the US). Also, the losses caused by catastrophes rose from $48.7 billions between 1950 and 1988 to $98 billions between 89 and 98, that is, an increase of 101.2%.
2. Consolidation between largest reinsurers raises difficulties.

The amount of significant disasters is increasing. From 1970 to 1988: 11 events during 19 years. From 1989 to 1995: 19 events during 7 ans. We estimate that we should expect one sinister larger than $1 ($3) billion per year (per 2.5 years) currently.

Cat bond (cont'd)

For a typical cat bond, we build a special purpose vehicule (SPV) that acts like a reinsurer by issuing debt on the financial markets, and by providing a reinsurance policy to the insurer. The policy pays the insurer in case of a loss that is larger than some threshold.

The cashflows paid to the debt holders are rerouted to pay finance the insurers in case of sinister; the interest payments are suspended, and/or a fraction of the capital is cut.

Cat bond (cont'd)

Initially (1996), cat bonds only covered one particular time of sinister and a particular region with an horizon of one year.

Today, they often cover multiple types of sinisters and multiple country with an horizon of multiple years.

There have been 7 transactions in 97 ($1 billion) and 8 transactions in 98 ($1 billion). The average number and amount of sinisters are increasing every year.

Investors consider cat bonds as a new class of assets that offers returns uncorrelated to other usual products such as stocks and bonds. Such assets help reduce the risk of a portfolio because of diversification.

Example - Atlantic bassin hurricane

The main factor for sinisters is the wind intensity. We can model the arrival of high velocity streams with a Poisson process.

We compute the average arrival time of an event with some minimum amplitude, and the probability that there is no event with a large amplitude during a given period (e.g., a year).

The average arrival time is approximately 3.6 years for winds at 125 mph or more, and the probability that there is no event during one year is 74% (i.e., there is 26% probability that such an event happens during the year). We also model the atmosphere pressures as low pressure environments increase the probability of significant storms.