The strictly 1/2-stable example

1 Direct approach: building a Lévy pure jump process on $\mathbb{R}$

Bert Fristedt provided key mathematical facts for this example.

A pure jump Lévy process X is a Lévy process such that X_t for any time t is the sum of all the jumps it has taken up to that time. We want to build such a process X. The particular one we want is called the increasing strictly 1/2-stable process (Section 3.7, page 93 in Cont and Tankov).

1.1 The setting

We begin by thinking about telephone calls.

Suppose we are receiving a Poisson stream of calls. We start at height zero and whenever the telephone rings we jump upward a unit amount. Then our height, as a function of time, is a Poisson process.

Let us generalize this. Suppose that we receive telephone calls randomly, but now each call gives us instructions: "Please jump an amount h". If we start at height zero and jump according to the instructions, our height as a function of time is a random process X.

To describe this random process more precisely, we categorize each call according to the size h of the jump which the call instructs us to make. If the call tells us to jump by an amount h, we will say the call is of type h. Fix an interval $[a, b]$ with either $0 < a < b$ or $a < b < 0$. Consider all calls of type h, where $a \leq h \leq b$, that we have received by time t. The number of these calls is assumed to be a Poisson process, which we might denote by $N_t^{[a,b]}$. Furthermore we assume that for disjoint intervals $[a, b]$ and $[c, d]$ the processes $\left(N_t^{[a,b]}\right)_{t\geq 0}$ and $\left(N_t^{[c,d]}\right)_{t\geq 0}$ are independent. But what are the rates of these Poisson processes?

We will be quite general. We will describe our rates using a measure ν on sets of jump sizes. ν will be the Lévy measure for the process X that we are going to construct. It is also referred to as the *intensity measure* for the process X.

ν is assumed finite on any set which omits an neighborhood of 0. That is, $\nu(A) < \infty$ for any Borel set A such that 0 is not in the closure of A. For the random structure of the calls that we have described, we assume that the rate for $N_t^{[a,b]} = \nu([a, b])$, and hence that the mean value of $N_t^{[a,b]} = \nu([a, b])\, t$.

We will also assume $\nu(\{0\}) = 0$, i.e. no calls of type 0, since a jump by an amount 0 causes no change.

(Note that $\nu(B)$ is also a mean value of an appropriate Poisson random measure on $[0, \infty) \times \mathbb{R}$.)

Intuitively, the process X we are going to construct is such that at time t, the value X_t is the height we have reached by starting at height zero and jumping according to the telephone calls we have received.

It is important to point out that we are allowing ν to become infinite near 0. Of course, this is dangerous, because if we get too many jump requests, even though they are very small, the sum of the jumps may not converge. So we will have to place some restriction on ν. But we don't need to require that $\nu(\mathbb{R}) < \infty$.

(If it does happen that $\nu(\mathbb{R}) < \infty$, we have the situation of the *compound Poisson process*.)

We expect that it will be necessary to assume that

$$\int_{[-1,1]} x^2 \nu(dx) < \infty, \tag{1}$$

so that ν will satisfy the requirements of any Lévy measure. However, for the type of process we want, we must actually assume a *stronger condition*

$$\int_{[-1,1]} |x| \, \nu(dx) < \infty, \tag{2}$$

As Cont and Tankov explain in Proposition 3.9 on page 86, this condition is required for a Lévy process to have paths of finite total variation. Since the process we want is an increasing process, it definitely has paths of finite total variation. (In general a Lévy process with paths of finite total variation is equal to a pure jump process plus a drift.)

1.2 The case of a finite intensity measure

Before we concentrate on our particular example, let's think a little more about how we would construct a jump process with a *finite* intensity process. As we noted earlier, this is the case of a *compound Poission process.* Once we have studied this case, we'll return to the construction of the example

To start with, let's consider a really simple case. Suppose we have two telephones, telephone 1 and telephone 2.

Telephone 1 receives calls from Philadelphia, and the calls form a Poisson stream with rate λ_1. The number of calls received up to time t (from some initial time) form a Poisson process N_t^1.

Telephone 2 receives calls from Boston, and the calls form a Poisson stream with rate λ_2. The number of calls received up to time t (from some initial time) form a Poisson process N_t^2.

We assume these two streams are independent, which may be a reasonable approximation for whatever situation we are modelling. A fundamental property of the Poisson distribution is that then the total number of class $N_t \equiv N_t^1 + N_t^2$ is also a Poisson process, with rate $\lambda = \lambda_1 + \lambda_2$.

Suppose that (for some reason), we can charge a client company \$2 for each call from Philly and \$3 for each call from Boston. We call these amounts the *values* of the calls. Then the total charges up to time t is a process $X_t \equiv 2N_t^1 + 3N_t^2$, and this is a typical compound Poisson process. The *intensity measure* ν for this process is $\lambda_1\delta_2 + \lambda_2\delta_3$. Notice we can check this from the definition, since for any subset B of the real line, the expected number of calls received in a unit time which have value in B is $\nu(B)$ (there are only 4 cases to check, depending on the number of elements in $\{2,3\} \cap B$).

Now suppose that we can't afford two telephones. We buy a single telephone, and route all the calls, both from Philly and from Boston, through our single telephone. The number of calls by time t is of course the process $N_t \equiv N_t^1 + N_t^2$ mentioned earlier. It is a Poisson process with rate $\lambda = \lambda_1 + \lambda_2$. But how can we charge for the calls now, when we don't know which calls are from which city? We can make a reasonable guess, as follows. For each call, we toss a coin which has success probability

$$p_1 = \frac{\lambda_1}{\lambda_1 + \lambda_2}$$

and failure probability

$$p_2 = 1 - p_1 = \frac{\lambda_2}{\lambda_1 + \lambda_2}.$$

If we obtain a success with our coin toss, we arbitrarily assign the call to Philly, and charge \$2. If we obtain a failure, we charge \$3.

It is a fundamental property of Poisson processes that this new way of billing gives us a process $\tilde{X}_t$ which is statistically indistinguishable from the previous process X, i.e. $\tilde{X}$ and X have the same distribution. Thus, though our records are faked, on the *average* we are charging the correct amount.

The lesson here is that there are two equivalent ways to construct a compound Poisson process with given intensity measure. We can either (i) add multiples of independent Poisson processes, or (ii) *split* a single Poisson process and multiply the pieces.

Method (i) seems more intuitive. Method (ii) seems to generalize more readily.

Let's now consider an arbitrary finite measure ν on the Borel subsets of the real line. We will construct a compound Poisson process which has ν as its intensity measure. We assume ν is nonzero. Define the probability measure f by

$$f(B) = \frac{\nu(B)}{\nu(\mathbb{R})}.$$

We imagine that we have a Poisson stream of telephone calls with rate $\lambda = \nu(\mathbb{R})$. The number of calls by time t is a Poisson process N_t with rate λ. When a call comes in, we proceed analogously to Method (ii) above and randomly assign the call a *value* to the call by sampling from the distribution f. The sum of the values for times up to time t is the compound process X with intensity measure ν. More precisely, we construct an iid sequence Z_n, $n = 1, 2, \ldots$, where each Z_n has distribution f. Then we let

$$X_t = \sum_{j=1}^{N_t} Z_j.$$

This is how Cont and Tankov define a compound Poisson process.

1.3 Approximating X

Now we return to the construction of our increasing example. We won't restrict ourselves to a particular Lévy measure in this part though.

Let C be a Borel subset of $\mathbb{R}$ such that 0 is not in the closure of C. Then $\nu(C) < \infty$. Suppose we decide to *ignore* calls that request jumps which are not in C. What kind of process do we get using the remaining calls?

The measure describing the calls that we accept will be denoted by ν^C, where

$$\nu^C(A) = \nu(A \cap C).$$

Define the probability measure f^C on $\mathbb{R}$ by

$$f^C(A) = \frac{\nu^C(A)}{\nu^C(\mathbb{R})}.$$

Let M be a Poisson process with rate $\nu^C(\mathbb{R})$. Let $(Z_n)_{n\geq 0}$ be an iid sequence such that each Z_n has distribution f^C. Define

$$X^C = \sum_{n=1}^{M} Z_n.$$

This compound Poisson process describes jumping according to the telephone calls which are of type h, where $h \in C$.

When C is a bounded set the integral of Z_n exists and is finite. In that case it is intuitively clear (and easy to prove) that

$$EX^C = \int_C x\nu(dx). \tag{3}$$

Now we would like to approximate X using the processes X^C. Since calls with small jumps are the problem for us, it's natural to start our construction of the true process X by fixing $\varepsilon > 0$ and taking $C = (-\varepsilon, \varepsilon)^c$. We can hope that X is the limit of the processes $X^{(-\varepsilon,\varepsilon)^c}$ as $\varepsilon \to 0$.

To simplify our discussion, let's focus on processes that only take positive jumps. That means that from now on we assume that

$$\nu((-\infty, 0]) = 0. \tag{4}$$

In this case $X^{(-\varepsilon,\varepsilon)^c} = X^{[\varepsilon,\infty)}$.

Consider a sequence $\varepsilon_k > 0$, $k = 0, 1, \ldots$ such that ε_k decreases to zero. We would like to show that $X^{[\varepsilon_k,\infty)}$ converges to a limit. We can always arrange our construction so that the random variables $X^{[\varepsilon_0,\infty)}$, and $X^{[\varepsilon_k,\varepsilon_{k-1})}$, $k = 1, 2, \ldots$ are independent, and set

$$X^{[\varepsilon_k,\infty)} = X^{[\varepsilon_0,\infty)} + \sum_{j=1}^{k} X^{[\varepsilon_k,\varepsilon_{k-1})},$$

since this definition gives the correct distribution. Hence we see that we would like to show that the infinite series

$$\sum_{k=1}^{\infty} X^{[\varepsilon_k,\varepsilon_{k-1})}$$

converges. The terms in the series are independent random variables, so if the series converges in the weakest sense (in distribution) then it converges almost surely. Furthermore, the Three Series Theorem of Kolmogorov applies.

Cont and Tankov discuss a similar situation on page 82 of Section 3.4. They point out that the Three Series Theorem shows why the fundamental condition (1) must hold for a Lévy measure, because the sum of the variances of the truncated series elements must be finite. However, the the Three Series Theorem also says that the sum of the means of the truncated series elements must converge.

By (3) we have

$$EX^{[\varepsilon_k,\varepsilon_{k-1})} = \int_{[\varepsilon_k,\varepsilon_{k-1})} x\nu\,(dx)\,.$$

In our case, because the random variables $X^{[\varepsilon_k,\varepsilon_{k-1})}$ are nonnegative, convergence of the sum of the means is easily seen to be equivalent to condition (2). Since we have assumed that $\nu\,((-\infty, 0]) = 0$, this is also equivalent to requiring that

$$\int_{[0,1]} x\nu\,(dx) < \infty. \tag{5}$$

Hence this is the requirement we will actually make for ν.

Remark Although the Three Series Theorem of Kolmogorov gives us insight in general, we don't really need it for our situation here. We have chosen our random variables $X^{[\varepsilon_k,\infty)}$ so that they are increasing. The Monotone Convergence Theorem then shows that $X^{[\varepsilon_k,\infty)}$ will converge to a finite limit if and only if (5) holds.

1.4 Choosing the Lévy measure

We have discussed the construction of a general increasing pure jump process. Now we specialize to the one we are actually interested in. Let ν have density

g, where

$$g(x) = \begin{cases} \frac{1}{\sqrt{2\pi}} x^{-3/2} & \text{for } x > 0, \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

According to Cont and Tankov, equation (3.31) on page 93, a strictly α-stable process X must satisfy the following scaling condition: for each $c > 0$,

$$\left(\frac{X_{ct}}{c^{1/\alpha}} \right)_{t \geq 0} \stackrel{d}{=} (X_t)_{t \geq 0} \,. \tag{7}$$

Let's verify that for our process X, with $\alpha = 1/2$.

First let's think about the Lévy measure for $(X_{ct})_{t \geq 0}$. Compared to the original process X, everything happens c times as fast for this new process. Hence the Lévy measure for $(X_{ct})_{t \geq 0}$ is exactly $c\nu$, and thus the density is cg.

Now consider the process $c^2 X$. Let μ denote the Lévy measure for this process, with density $\tilde{g}$. Compared to the original process X, the jumps of the new process are c^2 as large. Hence

$$\mu\left(c^2 B\right) = \nu(B)$$

for every Borel set B. Hence

$$\int_{c^2 B} \tilde{g}(y)\, dy = \int_B g(x)\, dx.$$

Letting $y = c^2 x$ we have

$$\int_B \tilde{g}\left(c^2 x\right)\, c^2 dx = \int_B g(x)\, dx.$$

Since this holds for all B we have

$$\tilde{g}\left(c^2 x\right)\, c^2 = g(x)$$

for all x. That is,

$$\tilde{g}(y) = g\left(c^{-2} y\right) c^{-2} = c^3 g(y)\, c^{-2} = cg(y)\,.$$

It follows that $c^2 X$ has the same Lévy measure as $(X_{ct})_{t \geq 0}$, so the processes are the same.

This shows that our constructed process X really is the strictly 1/2-stable process, as claimed.

Remark on uniqueness An *increasing* strictly α-stable process X is determined by α, up to a constant factor in the time scale. If we don't require that the process be increasing then there are more possibilities, since we could consider the difference of two independent increasing strictly α-stable processes.

1.5 Frequent jumps

Consider a time interval $[0, b]$, where $b > 0$ may be very short. Since $\nu((0,1]) = \infty$, we can find a sequence I_n of disjoint intervals which are subsets of $[0, 1]$, such that

$$\sum_{n=1}^{\infty} \nu(I_n) = \infty.$$

Let A_n be the event that X has a jump of size h during the time interval $[0, b]$, where $h \in I_n$. The events A_n are independent.

Since the number of calls during the time interval $[0, b]$ which have size in I_n is a Poisson random variable with mean $\nu(I_n)\, b$, we know that

$$P(A_n) = 1 - e^{-\nu(I_n) b}.$$

By calculus there exists a constant $c > 0$ such that for $0 \leq x \leq 1$ we have

$$1 - e^{-x} \geq cx.$$

Hence for all $x \geq 0$ we have

$$1 - e^{-x} \geq c\, x \wedge 1.$$

It follows that

$$\sum_{n=1}^{\infty} P(A_n) = \infty.$$

By the second Borel-Cantelli Lemma we know that with probability one X has a jump during $[0, b]$ with size in $[0, 1]$.

Thus X takes infinitely many jumps during every interval of time.

2 The connection with Brownian motion

Here we follow some material in N. Krylov's book (*Introduction to the Theory of Random Processes*, Theorem 2.6.1 and Exercise 5.2.11).

2.1 Defining the process via stopping times

Let $(W_t)_{t \geq 0}$ be the usual Brownian motion process on $\mathbb{R}$. For each $a \geq 0$, let

$$\tau_a = \inf\{t :\ t \geq 0,\ W_t \geq a\}.$$

Thus τ_a is the first hitting time of $[a, \infty)$ by W. Since W is a continuous process, it cannot reach (a, ∞) without first being equal to a. Thus τ_a is in fact the first time t that $W_t = a$.

It is intuitively reasonable from the stationary independent increments property of of Brownian motion that $(\tau_a)_{a \geq 0}$ is a process with with stationary independent increments. We will call this process X. We will show that X is the same increasing strictly α-stable process constructed already.

Since W is a continuous process, for any $a < b$ the process W cannot move instantly from a to b. It follows that X is a strictly increasing process.

It is easy to check from the definitions and the continuity of W that X is a caglad process. Naturally we could easily modify X to make it cadlag, if we had a reason to do so.

Consider a point a such that $\tau_a = s$ and $W_\cdot$ has a local maximum at time s. (There are *many* such points, because Brownian motion wiggles so much.) Let $\delta > 0$ be such that $W_t \le a$ for $t \in (s-\delta, s+\delta)$. It follows from the definitions that X has a jump at time a of size at least δ. We can show in this way that X has infinitely many jumps in every time interval. The set of jump times are dense but are countable since the jumps add to a finite sum.

2.2 The maximum process

Let

$$M_t = \max\{W_s : \ s \le t\}.$$

This process gives the maximum attained by Brownian motion up to time t, so it is necessarily nondecreasing and continuous. Notice that it has lots of flat spots, i.e. intervals of constancy. Indeed, using the scaling property of Brownian motion and the independent increments property, we can show that with probability one the complement of the union of the intervals of constancy is a Lebesgue null set.

The process $\tau_\cdot$ just constructed earlier satisfies the following equation, by simple logic:

$$\tau_a = \inf\{t : \ t \ge 0, \ M_t \ge a\}.$$

Actually $\tau_\cdot$ is is a kind of inverse to $M_\cdot$. Of course it is not the usual inverse, because $M_\cdot$ is not one-to-one. But we know how to find an "inverse" of a distribution function when we are constructing random variables, and the formula we use in that situation the same one that holds for $\tau_\cdot$, even though $M_\cdot$ is not a distribution function. So loosely speaking $\tau_\cdot$ is the inverse of $M_\cdot$.

Then *range* of $\tau_\cdot$ is uncountable, since $\tau_\cdot$ is strictly increasing and hence is one-to-one. One can show that the range of $\tau_\cdot$ has Hausdorff dimension 1/2.

2.3 The reflection principle

We would like to find the distribution of X_a. First we review a property of Brownian motion.

Lemma 1 (Reflection) *Let W be Brownian motion. Let $a, b \ge 0$. Then*

$$P\left(\max_{0\le t\le b} W_t < a\right) = P\left(|W_b| < a\right).$$

Proof Let $\tau = \tau_a \wedge b$.

Consider the map $\theta : C([0,\infty)) \to C([0,\infty))$ defined by

$$\theta(f)(t) = \begin{cases} f(t) & \text{for } t \le \tau_a, \\ f(\tau_a) - (f(t) - f(\tau_a)) & \text{for } t > \tau_a. \end{cases}$$

Since Brownian motion has no preference for left or right, we can convince ourselves that $\theta \circ W$ has the same distribution as W. Thus after time τ_a, W is equally likely to pursue the reflected version of its path as to pursue the original version of its path.

We note that since W_b has a density, $P(W_b = a) = 0$.

Combining these facts, and drawing the appropriate picture, we conclude that given $\tau_a \le b$, the probability that $W_b > a$ is exactly one-half. That is,

$$P\left(\max_{0\le t\le b} W_t \ge a\right) = 2P(W_b > a).$$

That is,

$$1 - P\left(\max_{0\le t\le b} W_t < a\right) = 1 - 2P(W_b > a) = P(|W_b| < a),$$

where we use the symmetry of W_b in the final inequality.

This proves the lemma.

3 The distribution of X_a

Using Lemma 1 we have

$$P(X_a > b) = P(|W_b| < a) = P\left(\sqrt{b}\,|W_1| < a\right) = 2\int_0^{a/\sqrt{b}} \frac{1}{\sqrt{2\pi}} e^{-x^2/2}\,dx,$$

where we rescaled the normal density to make the next step easier. The next step is in fact differentiation with respect to b. This shows that X_a has a density h_a on $(0,\infty)$ given by

$$h_a(x) = \frac{1}{\sqrt{2\pi}} a x^{-3/2} e^{-a^2/2x}. \tag{8}$$

It also follows from the fact that $P(X_a > b) = P(|W_b| < a)$ and the stationary increments property for X that X is *stochastically continuous*.

Hence X is a Lévy process. But which one? We will show below that X has the Lévy measure ν studied earlier and given by (6). Hence that X is the strictly 1/2-stable process. However, we can also check the strictly 1/2-stable property directly from the definition.

Fix $c > 0$. For $b > 0$ we have

$$
\begin{aligned}
P\left(X_{ca} > b\right) &= P\left(\max_{0 \le t \le b} W_t < ca\right) \\
&= P\left(\max_{0 \le t \le b} \frac{W_t}{c} < a\right) \\
&= P\left(\max_{0 \le t \le b} W_{t/c^2} < a\right) \\
&= P\left(\max_{0 \le s \le b/c^2} W_s < a\right) \\
&= P\left(X_a > b/c^2\right) = P\left(c^2 X_a > b\right).
\end{aligned}
$$

This shows that

$$
\left(X_{ca}\right)_{t \ge 0} \stackrel{d}{=} \left(c^2 X_t\right)_{t \ge 0}.
$$

Hence by definition X is strictly 1/2-stable.

3.1 Finding the Lévy measure directly

As part of the proof of the Lévy-Khinchin decomposition (see Krylov's textbook), one can prove that the Lévy measure measure ν (for *any* Lévy process X) has the following property.

$$
\int g \, d\nu = \lim_{n \to \infty} n E g\left(X_{1/n}\right), \tag{9}
$$

for any bounded continuous function with compact support whose support does not contain the origin.

Taking $a = 1/n$ in (8) shows easily that ν is given by (6).