Birth and Death

Robert Wolpert

April 25, 2016

1 Inhomogeneous Random Walks

Let's consider discrete time integer-valued time-homogeneous processes whose steps are of size at most one. These are characterized by the three probabilities

$$\begin{aligned} b_x &= \mathsf{P}[X_{t+1} = x+1 \mid X_t = x] \\ r_x &= \mathsf{P}[X_{t+1} = x \qquad \mid X_t = x] \\ d_x &= \mathsf{P}[X_{t+1} = x-1 \mid X_t = x], \end{aligned} \tag{1a}$$

mnemonic for "birth", "remain", and "death" respectively, which must sum to one. For states $a < x < b$ set

$$w(x) := \mathsf{P}_x[T_b < T_a]$$

where $T_c := \inf\{t \geq 0 : X_t = c\}$ for any state c (or $T_c = \infty$ if the state is never reached) and where P_x gives the probabilities of events for the process that begins at $X_0 = x$. Evidently $w(\cdot)$ satisfies the boundary conditions $w(a) = 0$ and $w(b) = 1$. For $x \in (a, b)$, looking ahead one step shows that it satisfies the equations

$$w(x) = b_x w(x+1) + r_x w(x) + d_x w(x-1)$$

or, since $b_x + r_x + d_x = 1$,

$$b_x[w(x+1) - w(x)] = d_x[w(x) - w(x-1)]$$

or

$$\begin{aligned} [w(x+1) - w(x)] &= \frac{d_x}{b_x}[w(x) - w(x-1)] \\ &= \prod_{a<y\leq x} \frac{d_y}{b_y} \quad [w(a+1) - w(a)] \\ &= (q_x/q_a) \quad [w(a+1) - w(a)] \end{aligned}$$

where

$$q_x := \prod_{0<y\le x} \frac{d_y}{b_y}.$$

Telescoping,

$$\begin{aligned}
1 = [w(b) - w(a)] &= \sum_{a\le y<b} [w(y+1) - w(y)] \\
&= \sum_{a\le y<b} (q_y/q_a) \quad [w(a+1) - w(a)], \text{ and} \\
\mathsf{P}_x[T_b < T_a] =: w(x) &= \sum_{a\le y<x} (q_y/q_a) \quad [w(a+1) - w(a)]. \text{ Dividing,} \\
&= \frac{\sum_{a\le y<x} (q_y/q_a)}{\sum_{a\le y<b} (q_y/q_a)} = \frac{\sum_{a\le y<x} q_y}{\sum_{a\le y<b} q_y}.
\end{aligned}$$

1.1 Example 1: Fair Game

If $d_x = b_x > 0$ for each x, then $q_x \equiv 1$ and $w(x) = (x-a)/(b-a)$. For example, successive fair \$1 bets beginning with a stake of 90 would reach \$100 before falling to zero with probability $90/100 = 0.9$.

1.2 Example 2: Roulette

In US casinos, a bet on "red" or "evens" pays off at even odds, but (because of the two green "0" and "00" outcomes, along with the numbers "1"–"36"), the probabilities of winning and losing are $b_x = 9/19$ and $d_x = 10/19$, respectively. Thus $q_x = (10/9)^x$ and

$$w(x) = \frac{\sum_{a\le y<x} (10/9)^{y-a}}{\sum_{a\le y<b} (10/9)^{y-a}} = \frac{1 - (10/9)^{x-a}}{1 - (10/9)^{x-b}}.$$

In the example above, with $a = 0$, $x = 90$, and $b = 100$ this is

$$= \frac{1 - (10/9)^{90}}{1 - (10/9)^{100}} \qquad \approx 0.349,$$

substantially below the 90% value for a fair game.

2 Stationary Processes

The process $\{X_t\}$ is transient if and only if

$$
\begin{aligned}
0 &< \lim_{b\to\infty} \mathsf{P}_x[T_b < T_a] \\
&= \lim_{b\to\infty} w(x) \\
&= \frac{\sum_{a\le y<x} q_y}{\sum_{a\le y<\infty} q_y}, \quad \textit{i.e.}, \text{ if} \\
\sum_{1\le y<\infty} q_y &= \sum_{1\le y<\infty} \frac{d_1 d_2 \cdots d_y}{b_1 b_2 \cdots b_y} < \infty. \qquad (1b)
\end{aligned}
$$

If transient, the process will visit each state at most finitely-many times, and may never visit states $y < x$. Conversely, if (1b) fails and $\sum_{y\ge 1} q_y = \infty$, the process is called "recurrent" and it will visit every state infinitely often.

If the process $\{X_t\}$ has a stationary distribution $\pi_x = \mathsf{P}[X_t = x]$ on $\mathbb{Z}_0 = \{0, 1, \dots\}$ then it must be recurrent, and moreover the transitions from any $x-1$ to x must occur exactly as often as transitions from x to $x-1$. This requires $\pi_x d_x = \pi_{x-1} b_{x-1}$, so

$$
\pi_x = \frac{b_{x-1}}{d_x}\, \pi_{x-1} = \prod_{0<y\le x} \frac{b_{y-1}}{d_y}\, \pi_0.
$$

Since $\sum_x \pi_x = 1$, we must have

$$
Z := \sum_{x=0}^{\infty} \frac{b_0 b_1 \cdots b_{x-1}}{d_1 d_2 \cdots d_x} < \infty, \text{ and} \qquad (1c)
$$

$$
\pi_x = Z^{-1} \prod_{0<y\le x} (b_{y-1}/d_y) \propto \frac{b_0 b_1 \cdots b_{x-1}}{d_1 d_2 \cdots d_x}. \qquad (1d)
$$

SO: the process $\{X_t\}$ described by (1a) is *transient* if and only if (1b) holds, is *positive recurrent* with stationary distribution (1d) if and only if (1c) holds, and is *null recurrent* (*i.e.*, returns to each state infinitely often, but after waiting times with infinite means) if Eqns (1b, 1c) both fail.

3 Continuous Time

Similarly, we can construct integer-valued stochastic processes indexed by continuous time that remain constant for exponentially-distributed waiting times at each state x before moving to $x \pm 1$ with specified probabilities. For any numbers λ_x^+ and λ_x^-, let the process remain at x for an interval of time with the $\mathsf{Ex}(\lambda_x)$ distribution where $\lambda_x := (\lambda_x^+ + \lambda_x^-)$, and then move as in (1a) with transition probabilities $b_x := \lambda_x^+/\lambda_x$, $r_x := 0$, and $d_x := \lambda_x^-/\lambda_x$. If we ensure the transitions don't occur so rapidly that the process "explodes" (a sufficient condition is that $\lambda_x \le a + bx$ for some $a > 0, b > 0$), then the resulting process is well-defined for all $t \ge 0$ and satisfies

$$\mathsf{P}[X_{t+\epsilon} = y \mid X_t = x] = o(\epsilon) + \begin{cases} \epsilon\lambda_x^+ & y = x+1 \\ 1 - \epsilon\lambda_x & y = x \\ \epsilon\lambda_x^- & y = x-1. \end{cases} \tag{2a}$$

If this process has a (summable) stationary distribution $\pi_x = \mathsf{P}[X_t = x]$, then the rates of transitions from $x \leftrightarrow x+1$ must be equal, *i.e.*, $\pi_x \lambda_x^+ = \pi_{x+1}\lambda_{x+1}^-$. From this we may identify the stationary distribution as $\pi_x = Z^{-1}\pi_x'$ where

$$\pi_x' := \prod_{0 \le y < x} \frac{\lambda_y^+}{\lambda_{y+1}^-}, \quad \text{provided that} \quad Z := \sum \pi_x' < \infty. \tag{2b}$$

3.1 Example 3: Imigration & Linear Death

For example, choose constants $i, d > 0$ and set $\lambda_x^+ := i$ and $\lambda_x^- := dx$, so

$$\mathsf{P}[X_{t+\epsilon} = y \mid X_t = x] = o(\epsilon) + \begin{cases} \epsilon i & y = x+1 \\ 1 - \epsilon(i + dx) & y = x \\ \epsilon dx & y = x-1 \end{cases} \tag{3a}$$

describing a population with constant imigration rate i, and death rate d. The stationary distribution satisfies $\pi_x \propto \pi_x'$ with (by (2b))

$$\pi_x' = \frac{\lambda_0^+ \lambda_1^+ \cdots \lambda_{x-1}^+}{\lambda_1^- \lambda_2^- \cdots \lambda_x^-} = \frac{(i)(i)\cdots(i)}{(d)(2d)\cdots(xd)} = \frac{(i/d)^x}{x!}.$$

The sum $Z := \sum_{x=0}^{\infty} \pi_x' = e^{i/d}$ is finite, so the process is positive recurrent with stationary distribution

$$\pi_x = Z^{-1}\pi_x' = \frac{(i/d)^x}{x!}\ e^{-i/d}. \tag{3b}$$

Thus $X_s \sim \mathsf{Po}(i/d)$ has the Poisson distribution with mean (i/d). For $s \le t$ the function $\phi_x(t) := \mathsf{E}[X_t \mid X_s = x]$ satisfies initial condition $\phi_x(s) = x$ and, by (3a), ordinary differential equation (ODE)

$$\begin{aligned} \phi_x(t+\epsilon) &= \phi_x(t) + \epsilon[i - d\phi_x(t)] + o(\epsilon), \text{ hence} \\ \phi_x'(t) &= [i - d\phi_x(t)] \\ \left\{e^{d(t-s)}\phi_x\right\}' &= e^{d(t-s)}\left\{i - d\phi_x(t) + d\phi_x(t)\right\} = ie^{d(t-s)} \\ \left\{e^{d(t-s)}\phi_x\right\} &= x + \int_s^t ie^{d(u-s)}\,du \\ &= (x - i/d) + (i/d)e^{d(t-s)}. \end{aligned}$$

The solution is

$$\phi_x(t) = (i/d) + (x - i/d)e^{-d(t-s)},$$

beginning at $\phi_x(s) = x$ and converging as $t \to \infty$ at rate d to the mean (i/d), and

$$\mathsf{E}[(X_t - i/d)(X_s - i/d) \mid X_s = x] = (x - i/d)^2 e^{-d|t-s|} \text{ or, unconditionally,}$$
$$\mathsf{Cov}(X_t, X_s) = (i/d)e^{-d|t-s|}. \tag{3c}$$

The mean is the population size at which the death rate just matches the immigration rate, and the autocorrelation is of $\mathsf{AR}(1)$ form $\exp(-\lambda|t-s|)$, with decay rate $\lambda = d$.

3.2 Example 4: Imigration w/Linear Birth & Death

Now choose constants $i, b, d > 0$ and set $\lambda_x^+ := i + bx$ and $\lambda_x^- := dx$. This "Branching Negative Binomial" or $\mathsf{bNB}(i, b, d)$ process describes a population with constant immigration rate i and linear birth and death rates bx and dx, respectively, for a population of size x, so

$$\mathsf{P}[X_{t+\epsilon} = y \mid X_t = x] = o(\epsilon) + \begin{cases} \epsilon(i + bx) & y = x + 1 \\ 1 - \epsilon(i + (b+d)x) & y = x \\ \epsilon dx & y = x - 1 \end{cases} \tag{4a}$$

and the stationary distribution, if one exists, must be proportional to

$$\begin{aligned} \pi'_x &= \frac{\lambda_0^+ \lambda_1^+ \cdots \lambda_{x-1}^+}{\lambda_1^- \lambda_2^- \cdots \lambda_x^-} = \frac{(i)(i+b)\cdots(i+b(x-1))}{(d)(2d)\cdots(xd)} \\ &= \frac{b^x \Gamma(x + i/b)/\Gamma(i/b)}{d^x \ x!} \\ &= \binom{i/b + x - 1}{x} (b/d)^x = \binom{-i/b}{x} (-b/d)^x. \end{aligned}$$

By the binomial theorem

$$Z := \sum_{x=0}^{\infty} \pi'_x = (1 - b/d)^{-i/b}$$

is finite, so the process *is* recurrent with stationary distribution

$$\pi_x = Z^{-1}\pi'_x = \binom{i/b + x - 1}{x} (1 - b/d)^{i/b} (b/d)^x \tag{4b}$$

and $X_t \sim \mathsf{NB}(\alpha = i/b,\ p = 1 - b/d)$ has the negative binomial distribution with shape parameter $\alpha = i/b$ and success probability $p = 1 - b/d$) (and failure probability $q = b/d$). By lhe same ODE argument as before, for $s < t$

$$\mathsf{E}[X_t \mid X_s = x] = \big(\frac{i}{d-b}\big) + \big(x - \frac{i}{d-b}\big)e^{-(d-b)|t-s|}$$

converges exponentially at rate $(d-b)$ as $t \to \infty$ to the mean. The autocovariance is

$$\mathsf{Cov}(X_s, X_t) = \frac{i\ d}{(d-b)^2}\ e^{-(d-b)|t-s|} \tag{4c}$$

and the autocorrelation again is of $\mathsf{AR}(1)$ form $\exp(-\lambda|t-s|)$, now with decay rate $\lambda = (d-b)$.

Interesting, and makes sense. The autocorrelation drops if death rate is much higher than birth rate. The shape parameter hinges on the ratio of immigration to birth, and the mean is the population size at which deaths just balance immigration and births.

3.3 Hitting Probabilities

For states $L \le x \le R$ let

$$T_{LR} := \inf\{t :\ X_t \ge R \text{ or } X_t \le L \mid X_0 = x\}$$

be the first time X_t leaves the interval (L, R), and let

$$\psi(x) := \mathsf{P}_x[X_t \ge R \text{ before } X_t \le L \mid X_0 = x]$$

be the probability that $X_t = R$ at the stopping time T_{LR}. For small $\epsilon > 0$,

$$\psi(x) = \psi(x)(1 - [i + (b+d)x]\epsilon) + \psi(x+1)[i + bx]\epsilon + \psi(x-1)[dx]\epsilon + o(\epsilon).$$

Subtracting $\psi(x)$ from both sides, dividing by ϵ, and taking the limit as $\epsilon \to 0$,

$$0 = [i + (b+d)x][\psi(x+1) - \psi(x)] - [dx][\psi(x) - \psi(x-1)],$$

so ψ satisfies the equation

$$[\psi(x+1) - \psi(x)] = \frac{dx}{i + bx}[\psi(x) - \psi(x-1)]$$

These equations plus some boundary conditions will determine the function $\psi(x)$ for all $x \in \mathbb{N}$, not only on the interval $[L, R]$. Setting $c_0 := [\psi(1) - \psi(0)]$ and iterating, we have

$$\begin{aligned}
\psi(x+1) - \psi(x) &= \left[\frac{dx}{i+bx}\right]\left[\frac{d(x-1)}{i+b(x-1)}\right]\cdots\left[\frac{d}{i+b}\right]\ [\psi(1) - \psi(0)] \\
&= (d/b)^x \frac{x!\ \Gamma(1 + i/b)}{\Gamma(1 + i/b + x)}\ c_0
\end{aligned}$$

Setting $c_1 := \lim_{x\to\infty} \psi(x)$ and telescoping,

$$\begin{aligned}
\psi(x) &= c_1 - \sum_{j=x}^{\infty}[\psi(j+1) - \psi(j)] \\
&= c_1 - \sum_{j=x}^{\infty}(d/b)^j \frac{j!\ \Gamma(1 + i/b)}{\Gamma(1 + i/b + j)}\ c_0 \\
&= c_1 - c_0 g(x), \qquad \text{where} \\
g(x) &= \frac{x!\ \Gamma(1 + i/b)}{\Gamma(1 + i/b + x)}\,{}_2F_1(1, 1 + x; 1 + x + i/b; d/b)(d/b)^x.
\end{aligned}$$

The values of the constants c_0, c_1 are determined by the boundary conditions $\psi(L) = 0$, $\psi(R) = 1$, leading to:

$$\psi(x) = \frac{g(x) - g(L)}{g(R) - g(L)}.$$

4 Non-stationary Processes

Dynamic Birth/Death processes much like those of Section (3) can also be constructed with increase and decrease rates $\lambda_x^+(t)$ and $\lambda_x^-(t)$ that depend on time. For example, if any or all of (α, p, λ) or, equivalently, of (i, b, d) are locally L_1 functions of *time*, then we can construct a dynamic bNB process, a well-defined integer-valued Markov process X_t that satisfies

$$\mathsf{P}[X_{t+\epsilon} = y \mid X_t = x] = \begin{cases} o(\epsilon) + \epsilon(i_t + x b_t) & y = x+1 \\ o(\epsilon) + 1 - \epsilon[i_t + x(b_t + d_t)] & y = x \\ o(\epsilon) + \epsilon(x d_t) & y = x-1 \\ o(\epsilon) & \text{otherwise} \end{cases}$$

which can be constructed as follows.

Let $\{\zeta_n\} \overset{\text{iid}}{\sim} \mathsf{Ex}(1)$ be standard exponentially-distributed random variables. Fix any initial value $x_0 \in \mathbb{N}_0$ and initial time $t_0 \in \mathbb{R}$. For $n \geq 0$, set:

$$t_{n+1} := \inf\left\{ t > t_n : \int_{t_n}^{t} [i_s + (b_s + d_s) x_n] \, ds > \zeta_n \right\}$$

$$X_t := x_n, \quad t_n \leq t < t_{n+1}$$

$$x_{n+1} = \begin{cases} x_n + 1 & \text{w/prob } \dfrac{i_{t_{n+1}} + b_{t_{n+1}} x_n}{i_{t_{n+1}} + (b_{t_{n+1}} + d_{t_{n+1}}) x_n} \\ x_n - 1 & \text{w/prob } \dfrac{d_{t_{n+1}} x_n}{i_{t_{n+1}} + (b_{t_{n+1}} + d_{t_{n+1}}) x_n}. \end{cases}$$

In the special case where i, b, and d are constant, $t_{n+1} = t_n + \zeta_n / [i + (b+d) x_n]$ and this reduces to the earlier construction. Even in the general case where the parameters aren't constant that's probably a close enough approximation for simulations, using b_t, d_t, and i_t evaluated at $t = t_n$ for drawing t_{n+1} and at $t = t_{n+1}$ for drawing x_{n+1}.

Mean Function

The conditional mean satisfies

$$\mathsf{E}[X_{t+\epsilon} \mid X_t = x] = x + \epsilon[i_t + (b_t - d_t) x] + o(\epsilon)$$

so the process will have a mean drift velocity of $[i_t + (b_t - d_t)X_t]$ at time t. The unconditional mean $\mu_t = \mathsf{E}[X_t]$ will satisfy the differential equation

$$\begin{aligned} \mu_{t+\epsilon} &= \mu_t + \epsilon[i_t + (b_t - d_t)\mu_t] + o(\epsilon), \ i.e. \\ \mu_t' &= i_t + (b_t - d_t)\mu_t \\ &= \lambda_t[\alpha_t/\beta_t - \mu_t], \end{aligned}$$

for NB shape parameter $\alpha_t := i_t/b_t$, success odds $\beta_t := p_t/(1-p_t) = (d_t - b_t)/b_t$, and autocorrelation rate $\lambda_t := (d_t - b_t)$. The solution starting at $\mu_t = x_s$ at time $t = s$ is

$$\mu_t = x_s e^{-\Lambda_s^t} + \int_s^t \lambda_u e^{-\Lambda_s^u} [\alpha_u/\beta_u] \, du$$

where $\Lambda_s^t := \int_s^t \lambda_u \, du$. Note the mean is no longer simply α_t/β_t as in the stationary case. During intervals where the parameters are constant the mean will converge exponentially at rate λ to α/β, but in intervals where the mean is increasing (resp., decreasing) then the mean μ_t will be less than (resp., greater than) α_t/β_t.

Rockfall Model

We chose the dynamic bNB process to model rockfall counts at the Soufrière Hills Volcano (SHV) on the Carribean island of Montserrat. Although several other processes are available with NB marginal distributions and $\mathsf{AR}(1)$-like covariance, the bNB process has the advantage that its likelihood function is available in closed form for irregularly-sampled observations X_{t_i}— see NB inference notes,

`http://www.stat.duke.edu/courses/Spring16/sta961/lec/NB-infer.pdf`.

Specifying this process requires the specification of three functions of time, which we took to be the mean μ_t, odds β_t, and autocorrelation rate λ_t. We took $\lambda_t \equiv \lambda$ to be constant, and μ_t to be observed, and then considered three possibilities for α_t and β_t: constant shape $\alpha_t \equiv \alpha$, in which case the rate is $\beta_t = \alpha/\mu_t$; constant rate $\beta_t \equiv \beta$, in which case the shape is $\alpha_t = \beta\mu_t$; and an intermediate case in which we took the product $\alpha_t\beta_t = c^2$ to be constant, with

$$\alpha_t = c(\mu_t)^{1/2} \qquad \beta_t = c(\mu_t)^{-1/2} = c^2/\alpha_t$$

This third choice proved a better fit to the data than the other two.

To build a bNB-like process whose mean is some specified non-constant function μ_t with derivative μ_t', with constant values for $\lambda_t := \lambda$ and $\alpha_t\beta_t := c^2$, we need to solve the equation

$$\mu_t' = \lambda_t[\alpha_t/\beta_t - \mu_t] = \lambda[\alpha_t^2/c^2 - \mu_t]$$

for α_t and β_t. Evidently the solution is

$$\alpha_t = c\left\{\mu_t + \mu_t'/\lambda\right\}^{1/2} \qquad \beta_t = c\left\{\mu_t + \mu_t'/\lambda\right\}^{-1/2} \qquad \lambda_t := \lambda$$

if $\mu'(t)+\lambda\mu_t \geq 0$ for all t (*i.e.*, if $\mu_t e^{\lambda t}$ is everywhere increasing). If this condition is satisfied, then

$$i_t = \lambda\mu_t + \mu_t' \qquad b_t = \sqrt{\lambda i_t}/c \qquad d_t = \lambda + \beta_t.$$

This reduces to the earlier construction if μ_t is constant (and $\mu_t' := 0$), but it's a little different for non-constant μ_t, since $\mu_t \neq \alpha_t/\beta_t$ whenever $\mu_t' \neq 0$.

Extensions

This model cannot accomodate a desired mean function $\mu(t)$ that occasionlly has steep enough drops that the condition "$\mu'(t) + \lambda\mu_t \geq 0$" fails. Extreme values of $|\mu'(t)|$ are associated with times where the process is changing very rapidly, and it is unrealistic to expect that the period-one autocorrelation $\exp(-\lambda_t)$ would be large there. One possible way of modifying the model to handle this is to keep $\alpha_t\beta_t$ constant but weaken the condition that λ_t take the constant value λ by allowing λ_t to increase in turbulent times— for example, set $\lambda_t := \max\big(\lambda, r|\mu_t'/\mu(t)|\big)$ for any fixed $r > 1$ (say, $r = 2$), and set α_t and β_t as before, so

$$\begin{array}{lll} \alpha_t = c\left\{\mu_t + \mu_t'/\lambda_t\right\}^{1/2} & \beta_t = c\left\{\mu_t + \mu_t'/\lambda_t\right\}^{-1/2} & \lambda_t = \max\big(\lambda, r|\mu_t'/\mu(t)|\big) \\ i_t = \lambda_t\alpha_t/\beta_t & b_t = \lambda_t/\beta_t & d_t = \lambda_t + \beta_t. \end{array}$$

This reduces to the construction above if $|\mu_t'| \leq \lambda\mu_t/r$, but otherwise drops the autocorrelation when necessary to permit rapid changes in the mean.