Diffusion theory

The mean and variance function for discrete processes

Let us consider the stochastic geometric growth model $N_{t+1} = \Lambda_t N_t$, or equivalently the random walk $X_{t+1} = X_t + S_t$, where the Λ_t and $S_t = \ln \Lambda_t$ are independent with the same distribution. With initial population size $N_0 = \exp(X_0)$ at time 0 we find simply

$$X_t = X_0 + S_0 + S_1 + \ldots + S_{t-1}.$$

Hence, by the central limit theorem the distribution of X_t for a given X_0 is approximately normal with mean $X_0 + \mu t$ and variance νt, where $\mu = \mathrm{E} R_t$ and $\nu = \mathrm{var}(S_t)$. It is well known that this approximation is remarkably good, even for moderate values of t, which means that the form of the distribution of the S_t has practically no effect on the process X_t, only the expectation μ and the variance ν. We obtain a more general class of models of the type $X_{t+1} = X_t + S_t$, by allowing the distribution of S_t to depend on X_t. For this process let us write $\mu(x) = \mathrm{E}(S_t | X_t = x)$ and $\nu(x) = \mathrm{var}(S_t | X_t = x)$. In accordance with our remarks on the simple model with constant μ and ν, where the properties of the process is practically determined by these two parameters, we should expect that the functions $\mu(x)$ and $\nu(x)$ contains most of the information of the behavior of the general process. As an illustration we consider three processes with the same mean and variance functions, but with rather different distributions of the S_t for given values of X_t. Let the models be of the discrete logistic type $\mathrm{E}(\ln N_{t+1} | N_t = n) = \ln N_t + r(1 - n/K)$ or equivalently $\mathrm{E}(\Delta X_t | N_t = n) = \mathrm{E}(\Delta X_t | X_t = \ln n) = r(1 - n/K)$ giving

$$\mu(x) = r(1 - e^x / K)$$

where $x = \ln n$. We assume that the variance is constant

$$\nu(x) = \sigma^2.$$

The model may now be written as

$$\Delta X = \mu(x) + \sqrt{\nu(x)} U$$

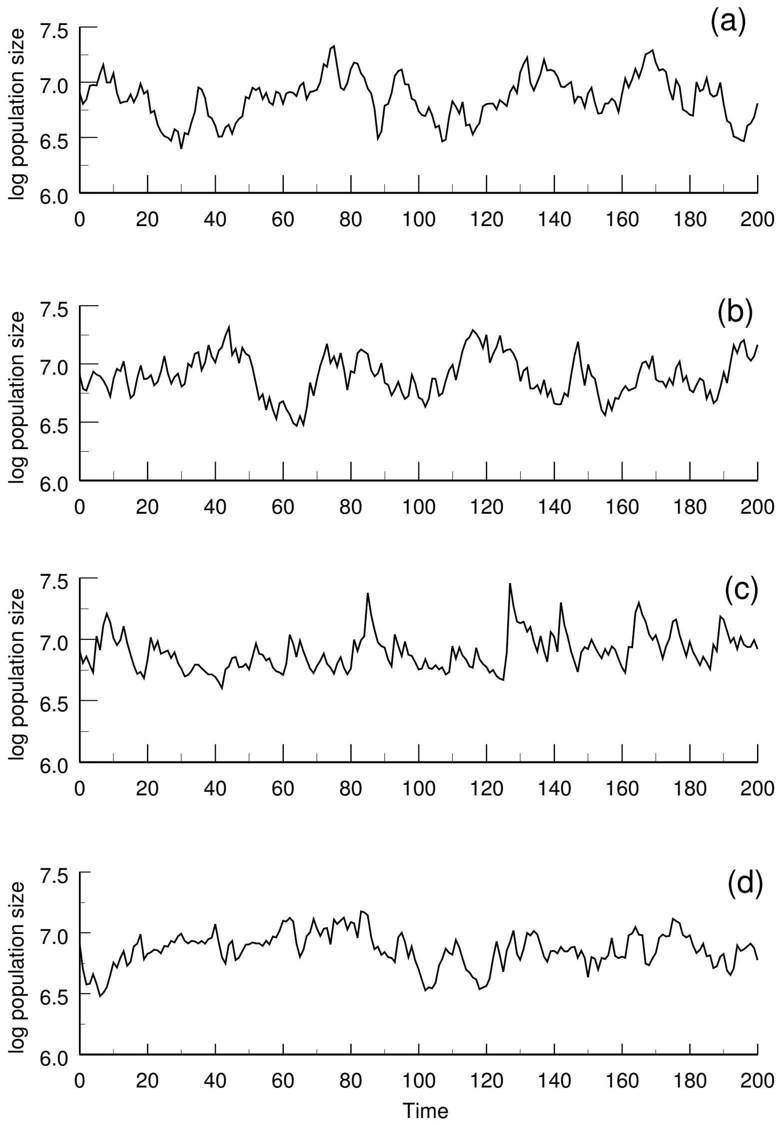

Figure 1: Population fluctuations for three models with the same discrete logistic type of dynamics with parameters $r = 0.2$, $K = 1000$, $\sigma^2 = 0.01$. The increments are modelled by different distributions: Normal distribution (a), Rectangular distribution (b), Exponential distribution (c) and the diffusion approximation recorded at discrete values with increments 1 (d).

where U is standardized so that $\mathrm{E}U = 0$ and $\mathrm{var}(U) = 1$.

Fig.1a-c shows simulations of this process when U is standardized normal, rectangular and exponential, respectively. Even if there is only one simulated process for each distribution we do get the impression that the fluctuations look fairly similar, especially for the normal and the rectangular, which both has zero skewness. For the exponential, which is skewed to the right, there is some tendency that the increases are somewhat quicker and the decreases somewhat slower than for the other distributions.

In the next section, we define diffusion processes, which is a class of processes that are continuous in the state variable as well as in time. The properties of such processes will be completely defined by the functions $\mu(x)$ and $\nu(x)$ which are called the *infinitesimal mean* and *infinitesimal variance* of the process. Together with possible boundary conditions, these functions completely define the diffusion process. It turns out that discrete processes often can be accurately approximated by diffusions with infinitesimal mean and variance equal to the mean and variance function of the discrete process.

The infinitesimal mean and variance of a diffusion

When the mean and variance functions are constants we have seen that the expectation as well as the variance of $X_t - X_0$ for a given X_0 are proportional to t, more precisely $\mathrm{E}(X_t - X_0|X_0) = \mu t$ and $\mathrm{var}(X_t - X_0|X_0) = \nu t$. The basic assumption of diffusions, apart from the Markov property the future depends only on the previous state), is that these relations hold for very small values of t, that is, for a small time interval Δt we assume $\mathrm{E}(\Delta X_t|X_t = x) \approx \mu(x)\Delta t$ and $\mathrm{var}(\Delta X_t|X_t = x) \approx \nu(x)\Delta t$. As Δt actually approaches zero we see that the last relation is equivalent to $\mathrm{E}[(\Delta X_t)^2|X_t = x] \approx \nu(x)\Delta t$ because $[\mathrm{E}(X_t|X_t = x)]^2$ is of order (Δt) and vanish compared to terms of order Δt as Δt approaches zero. The precise mathematical definitions are that the limit of $\mathrm{E}(\Delta X_t|X_t = x)/\Delta t$ as Δt approaches zero is the infinitesimal mean $\mu(x)$, while the limit of $\mathrm{E}[(\Delta X_t)^2|X_t = x]/\Delta t$ is the infinitesimal variance $\nu(x)$. Together with some boundary conditions, for example an extinction barrier, these two functions $\mu(x)$ and $\nu(x)$ completely define the diffusion process.

Diffusion processes may be simulated by using small discrete time steps. If the state at time t is $X_t = x$ we simulate

$$X_{t+\Delta t} = x + \mu(x)\Delta t + U_t\sqrt{\nu(x)\Delta t},$$

where the U_t are independent standard normal variables. By this method we obtain $\mathrm{E}(\Delta X_t|X_t = x) = \mu(x)\Delta t$ and $\mathrm{var}(\Delta X_t|X_t = x) = \nu(x)\Delta t$.

As an illustration Fig.1d shows one simulation of this process, that is $\mu(x) = r(1 - e^x/K)$ and $\nu(x) = \sigma^2$, serving as an approximation to all three processes shown in Fig.1a-c. The diffusion approximation constructed in this way, by choosing the mean and variance functions of the discrete process as the infinitesimal mean and variance, is commonly referred to as the Ito approximation. More precisely, the method is based on first expressing the process by a stochastic differential equation and using the stochastic integral called Ito integration when solving the equation, which is equivalent to dealing with the above diffusion.

Suppose now that we rather than working with $X_t = \log N_t$ considered the diffusion approximation to N_t. This diffusion approximation would then have infinitesimal mean and variance $\mu_N(n) = \mathrm{E}(\Delta N|N = n)$ and variance $\nu_N(n) = \mathrm{var}(\Delta N|N = n)$. It turns out that these two diffusions are not quite identical, but in practice fairly close if the changes in population size between years are not too large.

Boundary conditions

A diffusion is fully defined by its infinitesimal mean and variance together with some boundary conditions. In biology, the most actual boundary condition is defined by introducing an absorbing barrier at some value of N where the population actually goes extinct. Usually this extinction barrier is chosen at $N = 1$ or $N = 0$. When the population trajectory reaches the extinction barrier the population remains in this state. Hence, extinction barriers should only be used when modelling closed populations with no immigration from other populations.

Sometimes population models may also be defined by introducing a reflecting

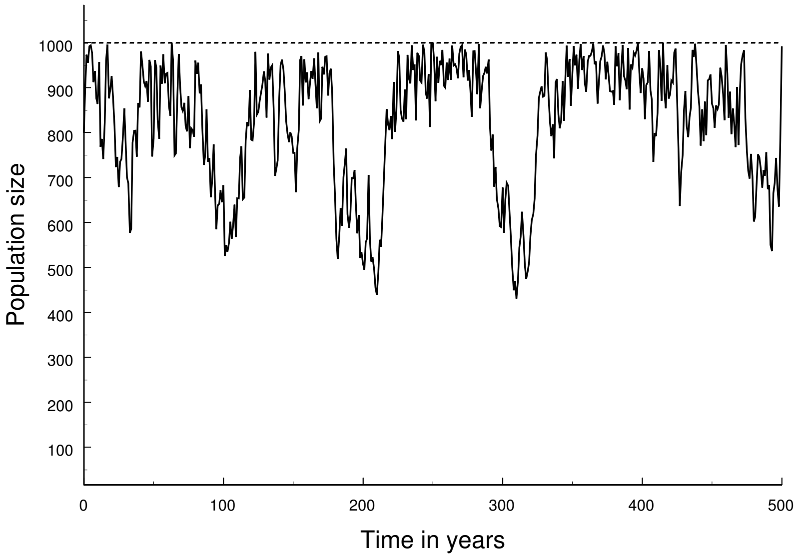

Figure 2: Simulation of the model with $\mu(n) = rn$ for $n < 1000$, a reflecting barrier at $n = 1000$ and infinitesimal variance $\nu(n) = \sigma^2 n^2$. The parameters are $r = 0.02$, $\sigma^2 = 0.01$.

barrier. This barrier can never be crossed. Rather than crossing, the process is immediately reflected. Mathematically, a reflecting barrier at, say $n = a$, can be modelled by defining the infinitesimal mean and variance to be symmetric around $n = a$. More precisely, for $n > a$ we use the infinitesimal mean $\mu(n) = \mu(2a - n)$ and variance $\nu(n) = \nu(2a - n)$ and treat the process as having no barrier at $n = a$. If the state of this process is $N_t > a$, we simply interpret this as if the state of the real process with reflecting barrier were $2a - N_t$.

Fig.2 shows a simulation of a model called 'geometric growth' with a reflecting barrier at population size 1000.

Some results valid for diffusion models

There are a number of useful analytical results on probabilities of absorbtion, distribution of time to absorbtion and stationary distributions for diffusions. Here we list some of these results.

The Wiener process

A process with constant environmental mean $\mu(x) = r$ and variance $\nu(x) = \sigma^2$ is called Wiener process. We consider such a process with initial state $x_0 > 0$ at time $t = 0$ and an absorbing barrier at zero.

If $r < 0$ this process will reach the absorbing barrier with probability 1. Absorbtion occurs at a time T which has the inverse Gaussian distribution

$$f(t) = \frac{x_0}{\sqrt{2\pi\sigma^2 t^3}} \exp[-\frac{(x_0 + rt)^2}{2\sigma^2 t}].$$

If $r < 0$ this is a proper distribution in the sense that the integral from zero to infinity of $g(t)$ is one. If $r \geq 0$, the process may be absorbed at infinity in which case the same integral equals the probability of ultimate extinction, which is $\exp(-2rx_0/\sigma^2)$. The cumulative distribution is given by

$$G(t) = P(T \leq t) = \Phi(-\frac{r_t + x_0}{\sigma\sqrt{t}}) + e^{-2rx_0/\sigma^2}\Phi(\frac{rt - x_0}{\sigma\sqrt{t}}),$$

where $\Phi(x) = \int_{-\infty}^{x} \exp(-x^2/2)dx$ is the standard normal integral.

If $r < 0$ the expectation is $\mathrm{E}T = -1/r$, and if $r > 0$ we have conditionally that $\mathrm{E}(T|T < \infty) = 1/r$.

In the case of no extinction barrier the distribution of X_t is simply the normal distribution with mean $x_0 + rt$ and variance $\sigma^2 t$. When there is an extinction barrier at $X_t = 0$ this distribution is no longer applicable. At a given time t, the process has either gone extinct, which occur with probability $G(t)$, or the population is still present with $X_t > 0$. The distribution of X_t in the case of an extinction barrier is also known. This probability density takes the form

$$h(x;t) = \frac{1}{\sqrt{2\pi t}\sigma}[1 - e^{-2xx_0/(\sigma^2 t)}]e^{-(x-x_0-rt)^2/(2\sigma^2 t)}.$$

This distribution may be integrated to give $P(X_t > x) = \int_x^\infty h(z;t)dz$, giving

$$H(x;t) = P(X_t \leq x) = \Phi(\frac{x - x_0 - rt}{\sigma\sqrt{t}}) + e^{-2rx_0/\sigma^2}\Phi(\frac{rt - x - x_0}{\sigma\sqrt{t}})].$$

Notice that $H(0;t) = P(X_t = 0) = P(T \leq t) = G(t)$ as expected.

Geometric Brownian motion

The process defined by $\mu(x) = rx$ and $\nu(x) = \sigma^2 x^2$ is called a geometric Brownian motion. One can show that the process $Y_t = \ln X_t$ then is a Wiener process with infinitesimal mean $r - \sigma^2/2$ and variance σ^2. Hence, starting at $x_0 = \exp(y_0)$ the above results for the time to absorbtion are valid replacing x_0 by $y_0 = \ln(x_0)$ and r by $r - \sigma^2/2$ and choosing 1 as the absorbing state for X_t (and 0 for Y_t).

Geometric growth model with variance proportional to x

Consider the model $\mu(x) = rx$ and $\nu(x) = \sigma^2 x$. Then the time to absorbtion at zero has the simple form

$$P(T < t) = \exp[-\frac{2x_0 r e^{rt}}{\sigma^2(e^{rt} - 1)}]$$

for $r \neq 0$, and $\exp[-2x_0/(\sigma^2 t)]$ for $r = 0$. For $r > 0$ the limit obtained as t approaches infinity gives that the probability of ultimate extinction at zero is $\exp(2rx_0/\sigma^2)$. For $r \leq 0$ we see from the cumulative distribution that the probability of ultimate extinction is a certain event.

The Ornstein-Uhlenbeck process

The process with linear infinitesimal mean $\mu(x) = \alpha - \beta x$, where $\beta > 0$, and constant infinitesimal variance $\nu(x) = \sigma^2$ is called an Ornstrein-Uhlenbeck process. one can show that the distribution of X_t is normal with mean

$$\mathrm{E}(X_t|X_0 = x_0) = \alpha/\beta + (x_0 - \alpha/\beta)e^{-\beta t}$$

$$\mathrm{var}(X_t|X_0 = x_0) = \frac{\sigma^2}{2\beta}(1 - e^{-2\beta t}).$$

We see that the expectation as well as the variance tend to a limit as t approaches infinity. Hence, if t is sufficiently large, there is no information left in the observation $X_0 = x_0$, and the process will go on fluctuating forever around the mean α/β with variance $\sigma^2/(2\beta)$. Such processes are called stationary processes, and the distribution of the state for large values of t, $f(x) = f_\infty(x; x_0)$, is called the stationary distribution of the process. Accordingly, the Ornstein-Uhlenbeck process has a stationary distribution which is normal with mean α/β and variance $\sigma^2/(2\beta)$. The joint distribution of X_t and X_s, for $0 < t < s = t + h$ is the binormal distribution with

$$\mathrm{cov}(X_t, X_{t+h}) = \frac{\sigma^2}{2\beta}(1 - e^{-2\beta t})e^{-\beta h}.$$

Hence, when the process has reached stationarity the autocorrelation function is simply $\mathrm{corr}(X_t, X_{t+h}) = e^{-\beta h}$.

Model Gompertz type of model

Models with infinitesimal mean $\mu(x) = rx(1 - \ln x/\ln K)$ are called models of the Gompertz type. If $\nu(x) = \sigma^2 x^2$ one can show that $Y_t = \ln X_t$ then is the Ornstrein-Uhlenbeck process with infinitesimal mean $r - \sigma^2/2 - y/\ln K$ and variance σ^2. Hence the stationary distribution of Y is normal and X must have the lognormal distribution. All results for the Ornstein-Uhlenbeck process can immediately be applied for the process Y and transformed back to find the corresponding result for X.

The logistic model

The model with $\mu(x) = rx(1 - x/K)$ is called the logistic model. If $\nu(x) = \sigma^2 x^2$ the stationary distribution is the gamma distribution

$$f(x) = \frac{\alpha^\beta}{\Gamma(\beta)} x^{\beta-1} e^{-\alpha x},$$

where the scale parameter $\alpha = 2r/(\sigma^2 K)$ and the shape parameter is $\beta = 2r/\sigma^2 - 1$. Using the well known expression for the mean and variance of the gamma distribution we find

$$\mathrm{E}X = \beta/\alpha = K(1 - \frac{\sigma^2}{2r})$$

$$\text{var}(X) = \beta/\alpha^2 = \frac{\sigma^2}{2r} K^2 (1 - \frac{\sigma^2}{2r}).$$

Model with stationary beta-distribution

A model often applied in genetical problems is $\mu(x) = ax - b(1 - x)$, $\nu(x) = \sigma_2 x(1 - x)$ where $0 < x < 1$. The stationary distribution for this model is the beta-distribution

$$f(x) = \frac{\Gamma(p+q)}{\Gamma(p)\Gamma(q)} x^{p-1} (1 - x)^{q-1}$$

where $p = 2a/\sigma^2$ and $b = 2b/\sigma^2$.

Model with stationary gamma distribution of the second kind

A model often used in describing fish populations is called the Beverton-Holt model. The model is usually used for age-structured populations, but in its simplest form the diffusion approximation has the form

$$\mu(x) = x \left(\frac{\alpha}{1 + \beta x} - \gamma \right)$$

and variance $\nu(x) = \sigma^2 x$. This process is stationary if $2\alpha - \sigma^2 - \gamma > 0$, and then stationary distribution is the beta distribution of the second kind

$$f(x) = \frac{\Gamma(p+q)\beta^p}{\Gamma(p)\Gamma(q)} \frac{x^{p-1}}{(1 + \beta x)^{p+q}}$$

where $p = 2(\alpha - \gamma)\sigma^2/2 - 1$ and $q = 1 + 2\gamma/\sigma^2$. The mean and variance of this distribution is $p/[\beta(q - 1)]$ and $p(p + q - 1)/[\beta^2 (q - 1)^2 (q - 2)]$. The variance is finite if $q > 2$.

The theta-logistic model

The theta-logistic model is a generalization of the logistic model with

$$\mu(x) = rx[1 - (x/K)^\theta]$$

$$\nu(x) = \sigma^2 x^2.$$

Writing $r = r_1/(1 - K^{-\theta})$ and keeping r_1 a positive constant this model is valid for all values of θ. The stationary distribution is the so-called generalized gamma distribution

$$f(x; K, \alpha, \theta) = \frac{|\theta|(\frac{\alpha+1}{\theta})^{\alpha/\theta}}{K\Gamma(\alpha/\theta)}(x/K)^{\alpha-1}e^{-\frac{(\alpha+1)}{\theta}(x/K)^{\theta}} \quad \text{for} \quad \theta \neq 0,$$

which has moments

$$\mathrm{E}X^p = \frac{K^p \Gamma(\frac{\beta-\theta+p\theta}{\theta^2})}{\Gamma(\frac{\beta-\theta}{\theta^2})(\frac{\beta}{\theta^2})^{p/\theta}}$$

for $p = 1, 2, \ldots,$ for $\theta > 0$, while for $\theta < 0$ the p'th moment exists for $p < 1 - \beta/\theta$. From this expression the mean and variance may be computed.

For $\theta = 1$ this is the gamma distribution with shape parameter α and scale parameter $(\alpha+1)/K$. Notice that K is a scale parameter of the distribution $f(x; K, \alpha, \theta)$, while α and θ are shape parameters. If $\theta = \alpha$ the distribution is Weibull and if $\theta = -1$ it is the inverse gamma.

To confirm that the limiting distribution as θ approaches zero is the lognormal we introduce $z = x/K$, and observe that the distribution of z is proportional to $z^{\alpha-1}e^{-(\frac{\alpha+1}{\theta})z^{\theta}}$. Expanding the exponent at $\theta = 0$ we find

$$\exp[-(\alpha+1)/\theta - 2\ln(z) - \frac{1}{2}(\alpha+1)\theta \ln(z)^2 + ...].$$

Absorbing the constant in the constant factor of the distribution and observing that $(\alpha+1)\theta$ approaches $\beta = \frac{2r_1}{\sigma_e^2 \ln K}$ as θ tends to zero, we find that the limiting distribution is proportional to $(1/z)\exp[-\ln(z) - \frac{1}{2}\beta\ln(z)^2]$ or proportional to $(1/z)\exp[-\frac{1}{2}\beta(\ln(z) + 1/\beta)^2]$. Hence, the limiting distribution of z is the lognormal distribution, and the corresponding distribution of $\ln(z)$ is $N[-1/\beta, 1/\beta]$. Finally, since $\ln x = \ln z + \ln K$ we see that the limiting distribution of $\ln(X_t)$ is normal with mean $\ln(K)[1 - \sigma^2/(2r_1)]$ and variance $\sigma^2 \ln(K)/(2r_1)$.

In the limit as θ approaches infinity we obtain the model with $\mu(x) = rx$ below K and a reflecting barrier at K. We observe that the stationary distribution for this model is

$$f(x; K, \alpha, \infty) = \frac{\alpha}{K^{\alpha}} x^{\alpha-1}$$

for $0 \leq x \leq K$, and otherwise zero.