

Introduction to SDE Models in Population Biology

Lecture 2: Population Survival, Extinction, and Persistence

2-hour PhD short course – 50-page beginner version

Learning objectives and two-hour plan

Students should learn to

- obtain global solutions using stopping times and Lyapunov estimates;
- prove positivity and invariance of biologically meaningful state spaces;
- derive extinction and persistence criteria for scalar population diffusions;
- compute the stationary density of the stochastic logistic equation;
- connect invariant distributions with long-time biological averages.

Suggested timing: 0–30 min local/global solutions and stopping times; 30–55 min Lyapunov methods; 55–65 min break; 65–95 min extinction/persistence; 95–115 min stationary density and ergodic interpretation; 115–120 min summary.

Textbook guide for this lecture

The main tools come from:

- X. Mao: local/global existence, stopping times, Lyapunov functions, stochastic stability;
- B. Øksendal: generators, diffusion processes, Itô/Dynkin formulas, stationary viewpoints.

We use the stochastic logistic equation throughout:

$$dX_t = X_t(a - bX_t)dt + \sigma X_t dW_t, \quad a, b, \sigma > 0,$$

as a complete example where positivity, extinction, persistence, and a stationary distribution can all be seen explicitly.

Local strong solutions and explosion time

For

$$dX_t = b(X_t)dt + \sigma(X_t)dW_t,$$

local Lipschitz continuity gives a unique strong solution up to an explosion time τ_e . A convenient localization is

$$\tau_n = \inf\{t \geq 0 : |X_t| \geq n\}.$$

Then

$$\tau_n \uparrow \tau_e.$$

To prove global existence, it is enough to show that for every $T > 0$,

$$\mathbb{P}(\tau_n \leq T) \longrightarrow 0.$$

This is the basic stopping-time strategy used repeatedly in nonlinear population SDEs.

Strong solution, weak solution, and pathwise uniqueness

A **strong solution** is constructed on a given probability space with a given Brownian motion and is adapted to its filtration. A **weak solution** allows the probability space and Brownian motion to be part of the construction. **Pathwise uniqueness** means that two solutions driven by the same Brownian motion with the same initial condition coincide almost surely. For this short course we mainly work with strong solutions, because population trajectories are interpreted as responses to a specified environmental noise path.

Why explosion is a separate issue

Local existence does not automatically imply a solution for all $t \geq 0$. Define stopping times

$$\tau_n = \inf\{t \geq 0 : |X_t| \geq n\}, \quad \tau_\infty = \lim_{n \rightarrow \infty} \tau_n.$$

If $\tau_\infty < \infty$, the solution becomes unbounded in finite time. To prove global existence one shows

$$\mathbb{P}(\tau_\infty = \infty) = 1.$$

Lyapunov functions provide a systematic method to do this even when coefficients have superlinear growth.

A standard Lyapunov non-explosion principle

Suppose there exists $V \in C^2(\mathbb{R}^d; [0, \infty))$ such that

$$V(x) \rightarrow \infty \quad \text{as } |x| \rightarrow \infty,$$

and for some constant C ,

$$LV(x) \leq C(1 + V(x)).$$

Then Itô's formula applied to $V(X_{t \wedge \tau_n})$ gives

$$\mathbb{E} V(X_{t \wedge \tau_n}) \leq V(X_0) + C \int_0^t (1 + \mathbb{E} V(X_{s \wedge \tau_n})) ds.$$

Gronwall's inequality yields a bound independent of n , which can be used to show $\tau_e = \infty$ almost surely.

Lyapunov non-explosion: a standard proof pattern

Suppose $V \geq 0$, $V(x) \rightarrow \infty$ as $|x| \rightarrow \infty$, and

$$LV(x) \leq C(1 + V(x)).$$

Apply Itô's formula to the stopped process $X_{t \wedge \tau_n}$:

$$\mathbb{E}V(X_{t \wedge \tau_n}) \leq V(X_0) + C \int_0^t (1 + \mathbb{E}V(X_{s \wedge \tau_n})) ds.$$

Gronwall's inequality gives a bound independent of n . Since V is large on $\{|X| = n\}$, Markov's inequality forces

$$\mathbb{P}(\tau_n \leq t) \rightarrow 0.$$

Therefore finite-time explosion is excluded.

Gronwall inequality in SDE estimates

A deterministic inequality repeatedly used in SDE analysis is: if

$$u(t) \leq A + B \int_0^t u(s) ds,$$

then

$$u(t) \leq Ae^{Bt}.$$

After taking expectations, stochastic integral terms often disappear, leaving an inequality of exactly this type. Thus many SDE estimates follow the chain

Itô formula $\rightarrow$ take expectations $\rightarrow$ Gronwall.

The infinitesimal generator

For

$$dX_t = b(X_t)dt + \Sigma(X_t)dW_t,$$

with $a(x) = \Sigma(x)\Sigma(x)^\top$, the generator is

$$LV(x) = \nabla V(x) \cdot b(x) + \frac{1}{2} \text{tr}(a(x)D^2 V(x)).$$

Itô's formula can be written as

$$V(X_t) = V(X_0) + \int_0^t LV(X_s)ds + M_t,$$

where M_t is a local martingale. **Interpretation:** LV is the expected instantaneous rate of change of the observable V .

Generator calculations: three useful test functions

For

$$dX_t = b(X_t)dt + \sigma(X_t)dW_t,$$

$$Lf = bf' + \frac{1}{2}\sigma^2 f''.$$

Useful choices include

$f(x)$	$Lf(x)$
x	$b(x)$
x^2	$2xb(x) + \sigma^2(x)$
$\log x$	$b(x)/x - \sigma^2(x)/(2x^2)$

The first two control moments; the logarithm detects per-capita growth near the extinction boundary.

Dynkin formula as an expectation identity

Under suitable regularity,

$$\mathbb{E}f(X_t) = f(X_0) + \mathbb{E} \int_0^t Lf(X_s) ds.$$

More generally, for a bounded stopping time τ ,

$$\mathbb{E}f(X_{t \wedge \tau}) = f(X_0) + \mathbb{E} \int_0^{t \wedge \tau} Lf(X_s) ds.$$

This stopped form is especially useful because it allows us to apply estimates before we have already proved global boundedness or positivity.

Choosing Lyapunov functions in population models

Common choices and their roles:

Choice	Typical purpose
$V(x) = x^p$	moment estimates and growth bounds
$V(x) = \log x$	per-capita growth and extinction rates
$V(x) = x - c \log x$	control both $x \rightarrow \infty$ and $x \downarrow 0$
$V(x) = \sum_i (x_i - c_i \log x_i)$	positive-cone systems and recurrence

The art is to choose V so that LV has a useful sign or a useful upper bound.

A Lyapunov function for positivity and boundedness

For positive populations a common choice is

$$V(x) = x - 1 - \log x, \quad x > 0.$$

Then $V(x) \geq 0$, and

$$V(x) \rightarrow \infty \quad \text{as } x \downarrow 0 \quad \text{or} \quad x \rightarrow \infty.$$

Thus the same function detects both approach to the extinction boundary and escape to infinity. For multidimensional systems one often uses

$$V(x) = \sum_{i=1}^n c_i (x_i - 1 - \log x_i), \quad c_i > 0.$$

Stopping times near zero and infinity

For a positive scalar process, define

$$\tau_n = \inf\{t \geq 0 : X_t \notin (1/n, n)\}.$$

To prove positivity and non-explosion, estimate the stopped process on $[0, t \wedge \tau_n]$ and show

$$\mathbb{P}(\tau_n \leq t) \rightarrow 0.$$

Because leaving $(1/n, n)$ means either approaching zero or becoming very large, a coercive Lyapunov function controls both possibilities simultaneously.

Positivity: stopping before the boundary

For a positive population variable, define

$$\tau_n = \inf\{t \geq 0 : X_t \notin (1/n, n)\}.$$

On $[0, \tau_n]$ the logarithm is legitimate. If one can obtain estimates independent of n for a function such as

$$V(x) = x - 1 - \log x \geq 0,$$

then boundary hitting and explosion can be excluded by letting $n \rightarrow \infty$. **Why** $x - 1 - \log x$?

$$V(x) \rightarrow \infty \quad \text{both as } x \downarrow 0 \text{ and as } x \rightarrow \infty.$$

So one Lyapunov function controls both biologically dangerous directions.

Worked positivity estimate for logistic growth

For

$$dX = X(a - bX)dt + \sigma X dW, \quad V(x) = x - 1 - \log x,$$

we have

$$V'(x) = 1 - \frac{1}{x}, \quad V''(x) = \frac{1}{x^2}.$$

Hence

$$\begin{aligned} LV(x) &= (x - 1)(a - bx) + \frac{1}{2}\sigma^2 \\ &= -bx^2 + (a + b)x - a + \frac{1}{2}\sigma^2. \end{aligned}$$

The right-hand side is bounded above on $(0, \infty)$ because the leading term is $-bx^2$. Therefore

$$LV(x) \leq C.$$

This estimate, combined with localization, proves that a solution starting from $X_0 > 0$ stays in $(0, \infty)$ for all finite times.

Alternative positivity proof using logarithmic coordinates

Suppose

$$dX_t = X_t f(X_t) dt + \sigma X_t dW_t, \quad X_0 > 0.$$

As long as $X_t > 0$,

$$d \log X_t = \left(f(X_t) - \frac{\sigma^2}{2} \right) dt + \sigma dW_t.$$

If the right-hand side remains finite on every finite time interval, then $\log X_t$ cannot reach $-\infty$ in finite time. Hence X_t cannot hit zero. This argument is intuitive, but a fully rigorous proof is usually written with stopping times before applying Itô's formula to $\log X_t$.

What multiplicative noise does at the boundary

For

$$dX_t = X_t f(X_t) dt + \sigma X_t dW_t,$$

both the drift and diffusion vanish at $X = 0$. Thus 0 is an absorbing state: once the population reaches zero, it stays there. Starting from $X_0 > 0$, many standard population SDEs never hit zero in finite time but may satisfy

$$X_t \rightarrow 0 \quad \text{as } t \rightarrow \infty.$$

This distinction between *finite-time hitting* and *asymptotic extinction* is important.

The logarithmic equation revisited

For the stochastic logistic model,

$$d \log X_t = \left(a - \frac{\sigma^2}{2} - bX_t \right) dt + \sigma dW_t.$$

Integrating from 0 to t ,

$$\frac{\log X_t}{t} = \frac{\log X_0}{t} + a - \frac{\sigma^2}{2} - b \frac{1}{t} \int_0^t X_s ds + \sigma \frac{W_t}{t}.$$

Since

$$\frac{W_t}{t} \rightarrow 0 \quad \text{a.s.,}$$

the long-time behavior is governed by a balance between the corrected low-density growth rate and the density-dependent term.

Strong law for Brownian motion

A basic long-time fact is

$$\frac{W_t}{t} \rightarrow 0 \quad \text{almost surely.}$$

Therefore in a logarithmic identity of the form

$$\frac{\log X_t}{t} = \frac{\log X_0}{t} + \frac{1}{t} \int_0^t g(X_s) ds + \sigma \frac{W_t}{t},$$

the Brownian term does not contribute to the limiting average growth rate. This simple fact underlies many extinction and persistence threshold calculations.

A more general martingale law of large numbers

If M_t is a continuous local martingale with quadratic variation satisfying

$$\langle M \rangle_t = O(t),$$

then under standard conditions

$$\frac{M_t}{t} \rightarrow 0 \quad \text{a.s.}$$

For example,

$$M_t = \int_0^t h(X_s) dW_s, \quad \langle M \rangle_t = \int_0^t h^2(X_s) ds.$$

This extends the Brownian strong law and is essential when diffusion coefficients depend on the state.

Extinction when noise overwhelms growth

Because $X_t \geq 0$,

$$\frac{1}{t} \log X_t \leq \frac{\log X_0}{t} + \left(a - \frac{\sigma^2}{2} \right) + \sigma \frac{W_t}{t}.$$

If

$$a - \frac{\sigma^2}{2} < 0,$$

then

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log X_t \leq a - \frac{\sigma^2}{2} < 0 \quad \text{a.s.}$$

Consequently

$$X_t \rightarrow 0 \quad \text{exponentially fast a.s.}$$

Biological conclusion: environmental variability can cause extinction even though the deterministic growth parameter a is positive.

Persistence in mean when corrected growth is positive

Suppose

$$\alpha := a - \frac{\sigma^2}{2} > 0.$$

The integrated log equation becomes

$$b \frac{1}{t} \int_0^t X_s ds = \alpha + \frac{\log X_0 - \log X_t}{t} + \sigma \frac{W_t}{t}.$$

If one proves

$$\frac{\log X_t}{t} \rightarrow 0 \quad \text{a.s.},$$

then

$$\boxed{\frac{1}{t} \int_0^t X_s ds \longrightarrow \frac{\alpha}{b} = \frac{a - \sigma^2/2}{b} \quad \text{a.s.}}$$

This is persistence in time average, not convergence of X_t to a fixed constant.

Extinction proof for the logistic diffusion

For

$$dX_t = X_t(a - bX_t)dt + \sigma X_t dW_t,$$

we have

$$\log X_t = \log X_0 + \left(a - \frac{\sigma^2}{2}\right)t - b \int_0^t X_s ds + \sigma W_t.$$

Since the integral term is nonpositive,

$$\frac{\log X_t}{t} \leq \frac{\log X_0}{t} + a - \frac{\sigma^2}{2} + \sigma \frac{W_t}{t}.$$

If $a - \sigma^2/2 < 0$, then

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log X_t < 0,$$

so $X_t \rightarrow 0$ exponentially fast almost surely.

Persistence in mean: the key identity

Rearrange the logarithmic equation:

$$\frac{b}{t} \int_0^t X_s ds = a - \frac{\sigma^2}{2} + \frac{\log X_0 - \log X_t}{t} + \sigma \frac{W_t}{t}.$$

If one can prove

$$\frac{\log X_t}{t} \rightarrow 0 \quad \text{a.s.,}$$

then

$$\boxed{\frac{1}{t} \int_0^t X_s ds \rightarrow \frac{a - \sigma^2/2}{b}.}$$

This is the stochastic analogue of convergence to the deterministic carrying capacity a/b .

Why $\log X_t/t \rightarrow 0$ in the persistent regime?

A complete proof uses growth bounds for the solution. The main idea is:

- the negative quadratic drift $-bX_t^2$ prevents sustained super-exponential growth;
- moment estimates give enough control on large excursions;
- positivity prevents finite-time boundary problems;
- martingale terms grow sublinearly in t after normalization.

Thus $\log X_t$ is of lower order than t in the recurrent/persistent regime. **Teaching point:** the formal time-average calculation is simple, but a rigorous proof requires the solution estimates developed earlier in the lecture.

Stationary density from the Fokker–Planck equation

For a one-dimensional diffusion

$$dX = \mu(X)dt + s(X)dW,$$

a formal stationary density is proportional to

$$p(x) \propto \frac{1}{s^2(x)} \exp \left(\int^x \frac{2\mu(y)}{s^2(y)} dy \right).$$

For logistic growth,

$$\mu(x) = x(a - bx), \quad s(x) = \sigma x.$$

Therefore

$$p(x) \propto x^{2a/\sigma^2 - 2} \exp \left(-\frac{2b}{\sigma^2} x \right), \quad x > 0.$$

This density is integrable near 0 exactly when $a > \sigma^2/2$.

General one-dimensional stationary density formula

For

$$dX_t = b(X_t)dt + \sigma(X_t)dW_t,$$

a formal stationary density satisfies

$$0 = -\frac{d}{dx}(b(x)p(x)) + \frac{1}{2}\frac{d^2}{dx^2}(\sigma^2(x)p(x)).$$

Under a zero-probability-flux condition,

$$p(x) \propto \frac{1}{\sigma^2(x)} \exp\left(\int^x \frac{2b(y)}{\sigma^2(y)} dy\right).$$

This formula is useful for scalar diffusions, but integrability and boundary behavior must still be checked.

Deriving the Gamma density for logistic growth

For

$$b(x) = x(a - bx), \quad \sigma(x) = \sigma x,$$

the stationary-density formula gives

$$p(x) \propto x^{-2} \exp \left(\int^x \frac{2(a - by)}{\sigma^2 y} dy \right).$$

Hence

$$p(x) \propto x^{\frac{2a}{\sigma^2} - 2} \exp \left(-\frac{2b}{\sigma^2} x \right), \quad x > 0.$$

This is a Gamma density precisely when

$$a > \frac{\sigma^2}{2}.$$

The invariant law is a Gamma distribution

When $a > \sigma^2/2$, write

$$k = \frac{2a}{\sigma^2} - 1 > 0, \quad r = \frac{2b}{\sigma^2}.$$

Then the stationary density has Gamma form

$$p(x) = \frac{r^k}{\Gamma(k)} x^{k-1} e^{-rx}, \quad x > 0.$$

Its mean is

$$\mathbb{E}_\pi[X] = \frac{k}{r} = \frac{a - \sigma^2/2}{b}.$$

This agrees exactly with the long-time average obtained from the logarithmic equation. **This agreement is an early example of ergodic theory in action.**

Mean and variance under the invariant Gamma law

Write the stationary law as Gamma(shape k , rate λ) with

$$k = \frac{2a}{\sigma^2} - 1, \quad \lambda = \frac{2b}{\sigma^2}.$$

Then

$$\mathbb{E}_\pi[X] = \frac{k}{\lambda} = \frac{a - \sigma^2/2}{b},$$

consistent with the persistence-in-mean calculation. Also,

$$\text{Var}_\pi(X) = \frac{k}{\lambda^2} = \frac{\sigma^2(2a - \sigma^2)}{4b^2}.$$

The stationary mean is below the deterministic equilibrium a/b .

Mode of the stationary distribution

If $k > 1$, equivalently $a > \sigma^2$, the Gamma density has positive mode

$$x_{\text{mode}} = \frac{k-1}{\lambda} = \frac{a - \sigma^2}{b}.$$

If

$$\frac{\sigma^2}{2} < a \leq \sigma^2,$$

the invariant density is integrable but is largest near zero. Thus a population may be stochastically persistent while still spending substantial time at very low abundance.

Deterministic equilibrium versus stochastic equilibrium

Deterministic logistic ODE

- equilibrium $K = a/b$;
- $X(t) \rightarrow K$ for $X_0 > 0$;
- one asymptotic population size.

Stochastic logistic SDE

- sample paths keep fluctuating;
- stationary law describes long-term occupancy;
- average abundance is $(a - \sigma^2/2)/b$ when persistent.

Noise both widens the long-term distribution and lowers its mean relative to a/b .

Biological vocabulary and mathematical formulations

Biological question	Mathematical formulation	Typical tool
Stay positive?	invariance / non-hitting of 0	stopping times, log functions
Blow up?	non-explosion	Lyapunov estimates
Die out?	$X_t \rightarrow 0$ or negative log-growth rate	Itô formula for $\log X_t$
Persist?	positive time average / tightness away from boundary	growth rates, Lyapunov functions
Long-term distribution?	invariant measure / ergodicity	generator, recurrence, Fokker-Planck

Four distinct notions of long-time behavior

For stochastic populations it is important not to conflate:

- ① **almost-sure extinction:** $X_t \rightarrow 0$ a.s.;
- ② **persistence in mean:** $t^{-1} \int_0^t X_s ds$ has a positive limit or positive liminf;
- ③ **stochastic permanence:** the population remains in a compact positive range with high probability for large t ;
- ④ **stationarity/ergodicity:** the law approaches, or time averages follow, an invariant distribution.

Different theorems establish different notions, and one should state precisely which notion is proved.

Parameter sensitivity in the logistic model

The corrected growth rate

$$r = a - \frac{\sigma^2}{2}$$

shows a nonlinear effect of noise intensity.

- Increasing a raises persistence.
- Increasing b lowers the long-run mean but does not change the extinction threshold.
- Increasing σ lowers the effective growth rate and increases stationary variability.

This separation of roles is useful when interpreting fitted models or designing numerical experiments.

Students work on one of the following for 10 minutes:

- 1 Show that

$$V(x) = x - 1 - \log x$$

is nonnegative and diverges as $x \downarrow 0$ or $x \rightarrow \infty$.

- 2 Reproduce the calculation of LV for the logistic SDE and find an explicit upper bound.
- 3 Starting from the formal stationary-density formula, derive the Gamma parameters k and r .
- 4 Explain in words why the stationary density fails to be integrable when $a \leq \sigma^2/2$.

The critical case $a = \sigma^2/2$

For the stochastic logistic equation, the strict inequalities

$$a < \frac{\sigma^2}{2} \quad \text{and} \quad a > \frac{\sigma^2}{2}$$

lead to clean extinction and persistence conclusions. At the critical value

$$a = \frac{\sigma^2}{2},$$

the leading logarithmic drift vanishes and the density-dependent term becomes decisive:

$$d \log X_t = -bX_t dt + \sigma dW_t.$$

This borderline regime typically requires a separate argument; it is a good example of why threshold equalities should not be inferred automatically from the strict cases.

Beginner review: why stopping times are introduced

A stopping time is a random time whose occurrence can be decided from information available up to that time. Typical examples in population models are

$$\tau_n = \inf\{t \geq 0 : X_t \leq 1/n \text{ or } X_t \geq n\}.$$

Before τ_n , the process stays in the compact interval $[1/n, n]$. This is useful because coefficients that are not globally Lipschitz often become bounded and Lipschitz on compact sets.

Proof strategy

Prove estimates up to $t \wedge \tau_n$, then let $n \rightarrow \infty$ after showing that $\tau_n \rightarrow \infty$ a.s.

A first example of a stopping-time estimate

Let $X_t > 0$ solve a one-dimensional population SDE and define

$$\tau_n = \inf\{t \geq 0 : X_t \notin (1/n, n)\}.$$

For $t < \tau_n$, the logarithm $\log X_t$ is well defined and bounded between

$$-\log n \leq \log X_t \leq \log n.$$

Hence Itô's formula can be safely applied to $\log X_{t \wedge \tau_n}$. After deriving an estimate independent of n , one can pass to the limit. This is the standard way to avoid illegal calculations at the boundary $X = 0$.

Why Lyapunov functions often blow up at the boundary

To prove positivity, we want to prevent the process from reaching 0. A common choice is

$$V(x) = x - 1 - \log x, \quad x > 0.$$

Then

$$V(x) \geq 0, \quad V(x) \rightarrow \infty \text{ as } x \downarrow 0,$$

and also $V(x) \rightarrow \infty$ as $x \rightarrow \infty$. Thus boundedness of $\mathbb{E}V(X_{t \wedge \tau_n})$ controls both very small and very large values of the population.

Intuition

A Lyapunov function is a mathematical alarm: it becomes large near the undesirable region.

Computing a generator slowly

For

$$dX_t = b(X_t)dt + \sigma(X_t)dW_t,$$

the generator is

$$LV(x) = b(x)V'(x) + \frac{1}{2}\sigma^2(x)V''(x).$$

Example: for logistic growth with multiplicative noise,

$$b(x) = x(a - cx), \quad \sigma(x) = \sigma x.$$

If $V(x) = \log x$, then

$$V'(x) = \frac{1}{x}, \quad V''(x) = -\frac{1}{x^2}.$$

Therefore

$$L \log x = a - cx - \frac{\sigma^2}{2}.$$

This is exactly the drift of $\log X_t$.

Dynkin formula as the expectation of Itô's formula

Itô's formula gives

$$V(X_t) = V(X_0) + \int_0^t LV(X_s) ds + M_t,$$

where M_t is a martingale with mean zero under suitable integrability conditions. Taking expectations yields Dynkin's formula:

$$\mathbb{E}V(X_t) = V(X_0) + \mathbb{E} \int_0^t LV(X_s) ds.$$

How it is used

If $LV(x) \leq C - KV(x)$ outside a compact set, Dynkin's formula gives an average bound on $V(X_t)$ and helps prove recurrence or tightness.

Reading extinction proofs: the three-line structure

For many one-species models, the extinction proof has the following structure.

- 1 Apply Itô's formula to $\log X_t$:

$$\log X_t = \log X_0 + \int_0^t q(X_s) ds + M_t.$$

- 2 Bound the drift from above:

$$q(X_s) \leq q(0) = a - \frac{\sigma^2}{2} < 0.$$

- 3 Use $M_t/t \rightarrow 0$ a.s. to obtain

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log X_t < 0.$$

Then $X_t \rightarrow 0$ exponentially fast almost surely.

Reading persistence proofs: what changes?

When

$$a - \frac{\sigma^2}{2} > 0,$$

we no longer expect extinction. For the logistic SDE,

$$d \log X_t = \left(a - \frac{\sigma^2}{2} - cX_t \right) dt + \sigma dW_t.$$

Divide by t :

$$\frac{\log X_t - \log X_0}{t} = a - \frac{\sigma^2}{2} - c \frac{1}{t} \int_0^t X_s ds + \sigma \frac{W_t}{t}.$$

If $\log X_t/t \rightarrow 0$ and $W_t/t \rightarrow 0$, then

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t X_s ds = \frac{a - \sigma^2/2}{c}.$$

This is persistence in time average.

Stationary density: the beginner derivation

For a one-dimensional diffusion

$$dX_t = b(X_t)dt + \sigma(X_t)dW_t,$$

a stationary density p formally satisfies zero probability flux:

$$b(x)p(x) - \frac{1}{2} \frac{d}{dx} (\sigma^2(x)p(x)) = 0.$$

Thus

$$\frac{d}{dx} (\sigma^2(x)p(x)) = 2b(x)p(x).$$

Solving this first-order equation gives

$$p(x) = \frac{C}{\sigma^2(x)} \exp \left(\int^x \frac{2b(y)}{\sigma^2(y)} dy \right).$$

This formula is one of the most useful explicit tools in one-dimensional SDEs.

Gamma distribution reminder

A Gamma distribution with shape $\alpha > 0$ and rate $\beta > 0$ has density

$$p(x) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x}, \quad x > 0.$$

Its mean and variance are

$$\mathbb{E}X = \frac{\alpha}{\beta}, \quad \text{Var}(X) = \frac{\alpha}{\beta^2}.$$

For the stochastic logistic equation, the invariant density has this form when

$$a - \frac{\sigma^2}{2} > 0.$$

Thus persistence can be described not only by time averages but also by a full stationary distribution.

Beginner checklist for long-time analysis

When studying a new one-species population SDE, ask:

- 1 Does the solution remain positive?
- 2 Can explosion be ruled out?
- 3 What is the SDE for $\log X_t$?
- 4 What is the corrected low-density growth rate?
- 5 Does the martingale term satisfy $M_t/t \rightarrow 0$?
- 6 Is there a stationary density or invariant measure?
- 7 Does the invariant distribution have finite mean and variance?

These questions organize most of the calculations in this lecture.

Exercises

- Complete the localized Lyapunov proof of global positivity for the logistic SDE.
- Prove exponential extinction when $a - \sigma^2/2 < 0$.
- Verify directly that the Gamma density solves the stationary Fokker–Planck equation.
- Compute the stationary variance and compare it for two values of σ .

Reading

- Mao: Lyapunov methods, non-explosion, stochastic stability.
- Øksendal: generator, diffusion processes, and invariant/stationary viewpoints.