Population Dynamics Forced by Stochastic Catastrophic Events

Christopher Spalding

October 15, 2015

1 Introduction

In the modern world, population extinctions are often seen in somewhat of a negative light. However, within an evolutionary context their action has been instrumental in the sculpting of the modern biosphere. Removing old forms clears the way for the emergence of new species with novel traits. Despite such importance, the mechanisms by which extinctions play out in nature are still mysterious. For example, the Earth offered up 5 mass extinction events within the past half-billion years, but it is unknown whether similar biodiversity would exist today were it not for these events. Perhaps the smaller-scale stochastic forcing of climate and competition would have sufficed.

Traditionally, environmental forcing has been considered in two broad ways. Commonly, a coloured noise term is added to the equations of population dynamics, which represents environmental forcing over a given autocorrelation time [1, 2, 3]. Though not a population model, it has been shown that cells may undergo transitions between two bistable states over a timescale that is minimised by an appropriate choice for the autocorrelation time of the noise [3]. Such a dependence on timescale is interesting when one considers that both the cell problem and the extinction problem reduce to a mean 'exit time' problem [4], where one asks what the mean time is for stochastic dynamics to take a system through a given point (zero individuals in the case of extinction).

Another way of looking at environmental stochasticity is to consider that most environmentally driven extinction occurs during particularly detrimental events, known as Catastrophes. The typical strategy for analysis of these is to suppose that catastrophes come along at a prescribed rate ν and instantaneously remove a given fraction of the population $1-p$ each time they do so [5]. This method may be related to the former method by noting that the autocorrelation time and $1/\nu$ are qualitatively similar. Accordingly, just as there is an autocorrelation timescale that minimizes transition times in cells, there may exist a frequency over which catastrophes are the most detrimental in terms of population extinctions - *a most catastrophic catastrophe.*

The importance of the correlation time and/or the time between catastrophic events is of great interest because it suggests that despite the forcing being intrinsically random, the randomness has more of an effect upon some timescales than others. To emphasise the significance of this fact, consider a deterministic environmental forcing such as the day night or seasonal cycles. Many species have adapted their life cycles to take advantage of these predictable timescales, be it the breeding of mammals in spring or the growth-mortality daily cycles of phytoplankton [6]. It is less obvious whether a stochastic forcing, with no deterministic periodicity, can select for a specific timescale within a population.

To investigate the influence of time-scale upon population extinctions we model biological populations that experience stochastic catastrophic events with a typical frequency. To illustrate the problem, consider a favourable environment that is nevertheless struck randomly by deleterious events during which the death rate is enhanced. If these events are extremely rare, but last a long time, such as mass extinction events, the mean time to extinction of a given population might simply be the time until the next bad event. Consequently, if the events become more frequent, the extinction time goes down. Eventually, we reach a point where the events are not long-lived enough to make extinctions likely each time an event occurs. The extinction time then begins to lengthen as populations are usually able to survive through at least one event, having to wait for several before extinction occurs. This argument suggests that there is a frequency of stochastic events that minimises the extinction time.

2 Modelling Stochastic Populations

Not only are populations typically influenced by a stochastic environment, their dynamics are intrinsically stochastic. Specifically, births and deaths occur with a given probability, with the probability dependent upon the number of individuals. There always exists a possibility that no individual will reproduce over the individuals' lifetimes and thus extinction may occur by chance. Such intrinsic randomness is known as **demographic** stochasticity, distinct form **environmental** stochasticity. We investigate a population subject to both forms of stochastism.

We set up the problem as follows. Suppose there is a probability $P_n(\tau)$ that a population possesses n members at a time τ. Upon advancing time forward by one unit, the probability at each n will change in one of two ways. Probability may enter from some other n, which in the case of births means a probability flow from $n-1$ to n and deaths come from $n+1$. Alternatively, births and deaths at level n remove probability from P_n. We may write this process succinctly as

$$P_n(\tau+1) = \sum_m \pi_{nm} P_m(\tau), \tag{1}$$

where π_{nm} is the probability that the population moves from m to n. Population models usually consider only nearest neighbours to exchange probability, such as might be expected from single birth and death events. However, we retain generality at first.

One of the terms in the sum multiples $P_n(\tau)$ and so we pick it out

$$P_n(\tau+1) = \pi_{nn} P_n(\tau) + \sum_{m \neq n} \pi_{nm} P_m(\tau). \tag{2}$$

The probability of staying at the same n (π_{nn}) is just 1 minus the probability of transferring to anything else, or

$$\pi_{nn} = 1 - \sum_{m \neq n} \pi_{nm}. \tag{3}$$

Accordingly, we arrive at the discrete-time master equation

$$P_n(\tau+1) - P_n(\tau) = \sum_{m\neq n} \Big(\pi_{nm} P_m(\tau) - \pi_{mn} P_n(\tau)\Big). \tag{4}$$

It is often more convenient to take the continuous-time limit of equation (4), but we must then recast the probabilities π_{mn} as *rates* of probability flow R_{mn}. Upon doing so, we obtain the Master Equation

$$\frac{dP_n(t)}{dt} = \sum_{m\neq n} \Big(R_{nm} P_m(t) - R_{mn} P_n(t)\Big), \tag{5}$$

which may be written in matrix-vector form as

$$\dot{\mathbf{p}}(t) = \mathbf{Q}\,\mathbf{p}(t), \tag{6}$$

where we have defined the transition matrix $\mathbf{Q}$ and the vector $\mathbf{p}$ each of the elements of which corresponding to the probability at a given population level.

2.1 Birth-death model

Having set-up the Master Equation, we now prescribe forms for the elements of the transition matrix $\mathbf{Q}$. We require a death rate, δ_n and birth rate β_n as a function of individual number. Probability leaves step n if a birth or a death occurs, such that $R_{n+1,n} P_n = \delta_n$ and $R_{n-1,n} P_n = \beta_n$. Probability is gained by way of deaths from $n+1$ and births form $n-1$ such that the Master equation becomes

$$\frac{dP_n(t)}{dt} = -(\beta_n + \delta_n) P_n + \beta_{n-1} P_{n-1} + \delta_{n+1} P_{n+1}. \tag{7}$$

We will choose the exact forms for β_n and δ_n below. For now, suppose that there exists a stochastic, environmental variable $\mathcal{I}$ that modulates the death rate according to

$$\delta_n = \delta_n^{(0)} + \mathcal{I}(t) n, \tag{8}$$

where $\delta_n^{(0)}$ is the death rate when $\mathcal{I} = 0$.

Despite being probabilities, we have now introduced stochastic functions into β and δ themselves, removing the usefulness of the Master Equation. Strictly speaking, we must instead introduce a second stochastic dimension, described by a co-ordinate $\mathcal{I}$. However, for simplicity, we choose to allow $\mathcal{I}$ to take one of only two values, corresponding to a "good" state and a "bad" state. The good state has $\mathcal{I} = 0$ whereas in the bad state, $\mathcal{I} = aA$, where $A > 0$ is a constant and a is a constant with dimensions of inverse time that scales the organisms' generational overturning rate (see below).

The benefit of our two state system is that we can get around the requirement for a second dimension by defining two separate probability distributions, P_n^+ and P_n^-, where the former corresponds to the bad state and the latter the good state (the "plus" represents an enhanced death rate). We suppose that the system switches from the bad to the good state randomly, but with a typical frequency α. Conversely, the system switches from good to bad at rate $\epsilon\alpha$ where $\epsilon < 1$ corresponds to an environment that is in the good state more often than the bad state on average. [1] Note that the mean value of $\mathcal{I}(t)$, which we refer to as $\bar{A}$ is given by

$$\bar{A} = \frac{\epsilon}{1+\epsilon} A. \tag{9}$$

We are now in a position to write separate master equations for P_n^+ and P_n^-. In particular, we add a probability flow of αP_n^+ from the plus state to the minus state and likewise a flow of $\epsilon\alpha P_n^-$ from the minus state to the plus state. Including these terms, we obtain the Master Equations:

$$\begin{aligned}\frac{dP_n^+(t)}{dt} &= -(\beta_n + \delta_n^{(0)} + aAn)P_n^+ + \beta_{n-1}P_{n-1}^+ + \left(\delta_{n+1}^{(0)} + aA(n+1)\right)P_{n+1}^+ - \alpha\, P_n^+ + \epsilon\,\alpha\, P_n^- \\ \frac{dP_n^-(t)}{dt} &= -(\beta_n + \delta_n^{(0)})P_n^- + \beta_{n-1}P_{n-1}^- + \delta_{n+1}^{(0)}P_{n+1}^- + \alpha\, P_n^+ - \epsilon\,\alpha\, P_n^-. \end{aligned} \tag{10}$$

These equations describe the probabilistic trajectory of a population subject to the stochastic forcing described above. An alternative view points to consider trajectories of the population size. We do not adopt such an approach, but a typical realisation is illustrated in figure 6, where we plot the number of individuals (scaled by the carrying capacity) as a function of time under the influence of stochastic catastrophic events.

2.2 Matrix approach

It is more convenient to analyse the equations above in matrix form. Accordingly, we recast the Master Equations into the form:

$$\frac{d}{dt}\begin{pmatrix} P_n^+ \\ P_n^- \end{pmatrix} = \begin{pmatrix} M_{mn}^{(0)} + aAO_{mn} - \alpha I_{mn} & \epsilon\alpha I_{mn} \\ \alpha I_{mn} & M_{mn}^{(0)} - \alpha\epsilon I_{mn} \end{pmatrix} \begin{pmatrix} P_n^+ \\ P_n^- \end{pmatrix}, \tag{11}$$

where $M_{mn}^{(0)}$ is a tridiagonal matrix that represents birth and death rates in the good state. To include the environmentally-enhanced death rate, we define the matrix

$$O_{mn}P_n \equiv (n+1)P_{n+1} - nP_n, \tag{12}$$

along with the identity matrix, given by I_{mn}.

In order to complete the specification of the problem, we prescribe functional forms for the birth and death rates. We choose to have the birth rate grow linearly with n while the death rate grows quadratically, forcing an "equilibrium" number $n = K$ (the carrying capacity) at which $\delta_n^{(0)} = \beta_n$. The carrying capacity here corresponds to the stationary solution of the good state in the limit where dynamics are deterministic. Specifically, we define

$$\begin{aligned}\delta_n^{(0)} &= an\left(\frac{n}{K}\right) \\ \beta_n &= an \end{aligned} \tag{13}$$

---

[1] Switching between states in this manner is known as a Telegraph Process [4].

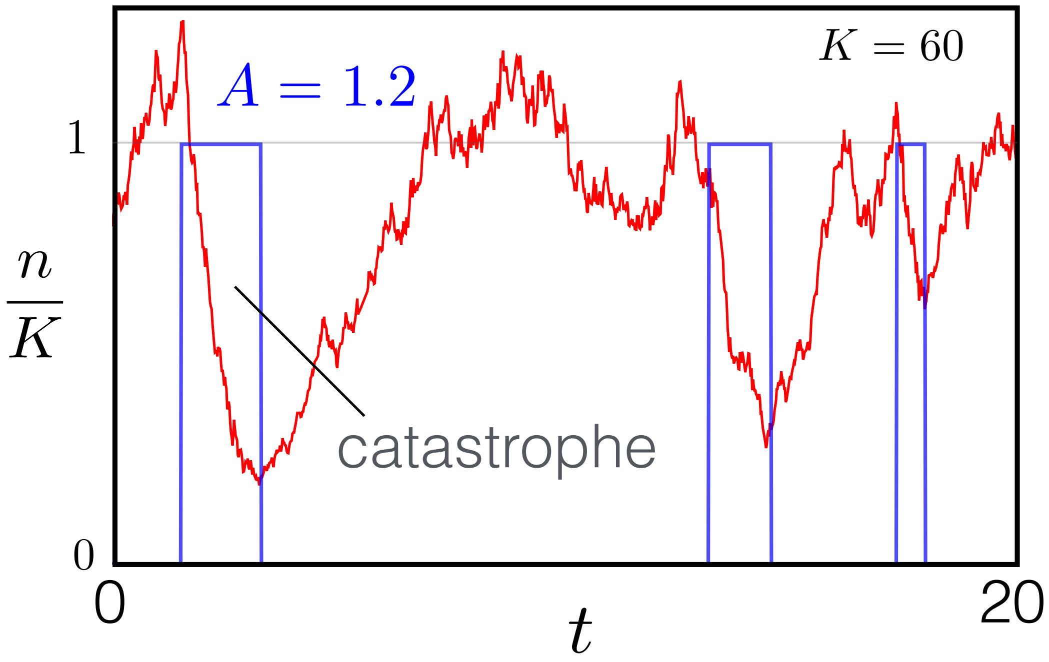

Figure 1: A typical realisation of the population size, n, normalised by the carrying capacity $K = 60$ under the action of stochastically distributed catastrophic events. The events typically last long enough to initiate a significant decline in the population numbers. The typical switching frequency $\alpha = 1$, with time measured in units of $1/a$. More rapid fluctuations would increase the extinction time because each catastrophe becomes significantly less detrimental. On the other hand, a slowing of the events would decrease extinction time, until the events' severity is offset by their scarcity.

and so as before, a measures the typical rate at which births and deaths proceed. Note that in this form, the linear environmental augmentation of death rate (An) naively appears commensurate with a decrease in birth rates. This statement is only true in the deterministic case; as we show below, adding both deaths and births generates more noise in the system than simply subtracting births, even if the resulting "deterministic" growth rate is the same in both cases.

2.3 Numerical solution

With all parts of the problem defined, we can now integrate equation (11) numerically to describe the time evolution of P_n under our specified stochastic forcing. The most important quantity to be extracted from the model is the extinction rate $\mathcal{R}_\mathrm{e}$ which for now we simply define as

$$\mathcal{R}_\mathrm{e} \equiv \frac{1}{P_0}\frac{dP_0}{dt}. \tag{14}$$

We present the time evolution of P_0 in Figure (3) along with its rate of change, using the initial condition that $P_n^+ = 0$ and P_n^- is drawn from the quasi-stationary distribution that would describe P_n in the absence of environmental forcing. This choice leads to a transient period, during which, probability flows from the minus state to the plus state until a quasi-steady state is reached for both. The timescale of the transient dynamics is related to how quickly equilibrium is established in the bad state.

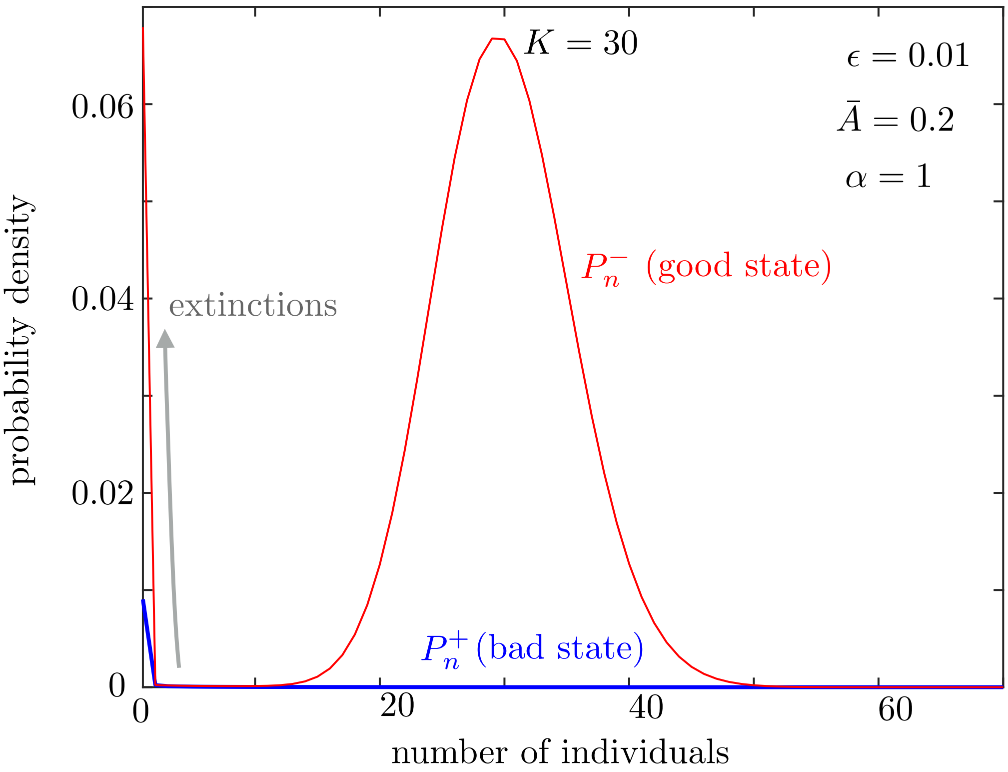

Figure 2: The quasi-steady distribution for the plus state (P_n^+, blue) and the minus state (P_n^-, red) resulting from an environmental perturbation with frequency α. Notice the pile-up of probability in the extinct state and the peak at $n = K$, but only for the minus state. The plus state only has a maximum at the origin.

In many problems, the transient period, which depends upon initial conditions, is of great importance. For example, initial conditions are crucial in computing the probability that any given event will lead to extinction. Additionally, if a system possesses multiple (quasi-) steady states, the initial configuration may determine where the long-time state will decay to. These complications do not apply to our considerations here, where we are interested in the long-term mean extinction time appropriate to an ensemble of populations under the influence of the environmental forcing prescribed above. Essentially, we neglect the probability that extinction occurs before the transient evolution decays.

Repeating the numerical integrations for a wide parameter-space would be time consuming. Instead, we take advantage of the quasi-steady evolution by seeking the lowest (in

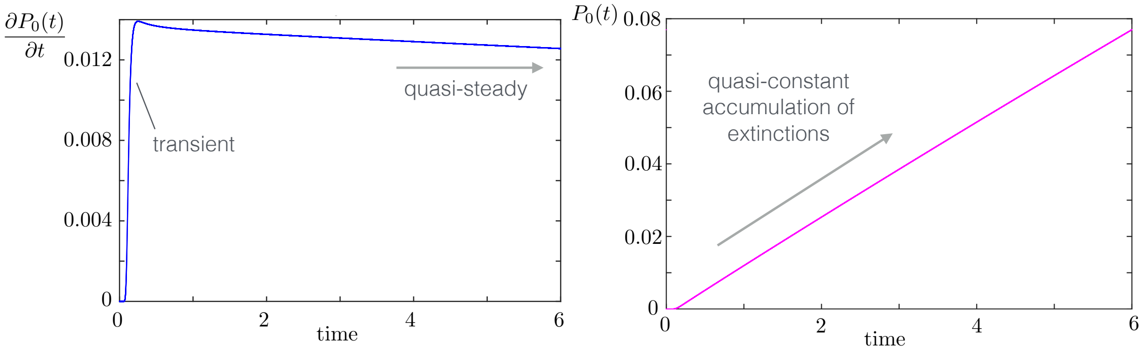

Figure 3: The flow of probability into the extinct state for the parameters of Figure 2. On the left is the rate of extinctions where the right indicates the probability of being extinct at a given time. After a brief transient period, extinctions begin to occur at a fairly constant rate. It is this quasi-steady rate we seek.

magnitude) eigenvalue λ_0 of the matrix on the RHS of equation (11). The inverse of λ_0 gives us approximately the mean time to extinction (MTE) of the population as we show next.

2.3.1 Mean time to extinction

After a time t, the probability a population is not extinct $\mathcal{P}(T_e > t)$ is simply the sum over all $n > 0$ of P_n,

$$\mathcal{P}(T_e > t) = \sum_{n>0} P_n(t), \tag{15}$$

and the probability distribution for extinction times is simply the (negative) time derivative of $\mathcal{P}(T_e > t)$. Suppose now that each P_n may be written as a sum of eigenmodes

$$P_n = \sum_m P_{nm} e^{-\lambda_m t}, \tag{16}$$

such that the mean extinction time can be written as

$$\begin{aligned} \bar{T}_e &= -\sum_{n>0}\sum_m P_{mn} \int_0^\infty T\left[\frac{\partial e^{-\lambda_m T}}{\partial T}\right] dT \\ &= -\sum_{n>0}\sum_m P_{mn}\frac{1}{\lambda_m}, \end{aligned} \tag{17}$$

where the second equality only holds when all $\lambda_m > 0$, i.e., there is no truly stationary solution except certain extinction ($P_0|_{t=\infty} = 1$). The transient state is rapid compared to the long-term quasi-steady decay of probability and so we may say that $\lambda_0 \ll \lambda_1 < \lambda_2$, etc. In other words, the sum over m above collapses to a single term, that of $m = 0$. Furthermore, the eigenvectors P_{mn} are unit-normalized, such that the double summation

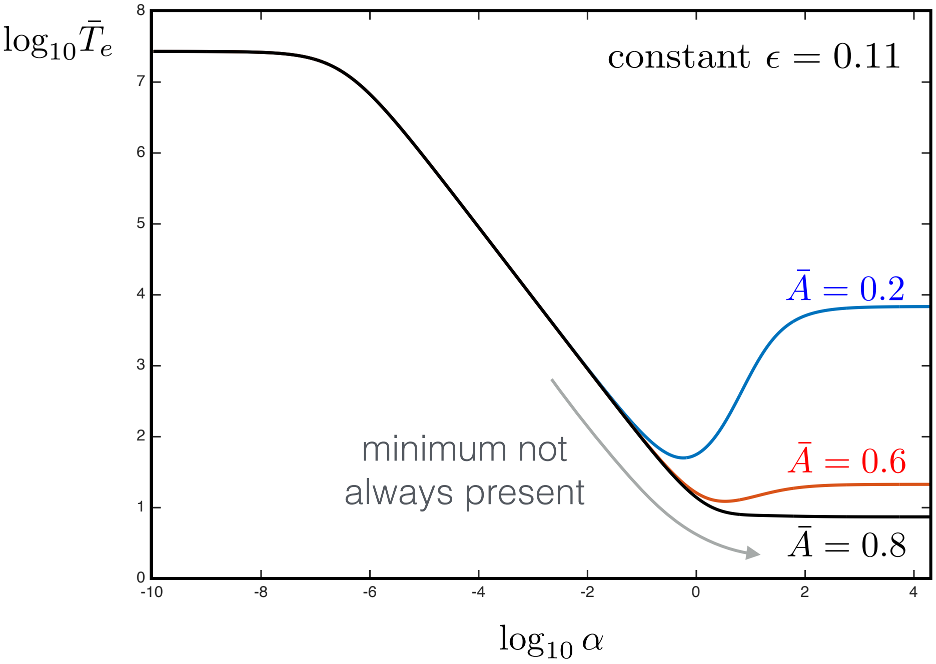

Figure 4: Similar to Figure 4, except we demonstrate that for large enough $\bar{A}$, the minimum ceases to exist.

reduces to a single inverse eigenvalue and the mean extinction is given by

$$\boxed{\bar{T}_e \approx \lambda_0^{-1}} \tag{18}$$

which is equivalent to our ascertain above.

2.4 Eigenvalue

We now compute the lowest eigenvalue numerically for a range of parameters. If $\bar{A} \gtrsim 1$, we find that the mean extinction time decreases monotonically with α (Figure 4). However, this is not the case we are interested in because such detrimental mean values make it unlikely a population would exist in the first place and, furthermore, the transient solution would become of significance in that case. Below, we consider only $\bar{A} < 1$, where the mean growth rate is not always negative, leaving room for stochastic periods of recovery. We present the results in Figure 5.

We immediately point out that there exists an α corresponding to a minimum in the mean extinction time. Such a minimum is reminiscent of several important results from the literature. For example, from the field of cell biology, cell-differentiation is sometimes thought to proceed by way of a transition from one stable state to another [7]. Theoretical models have shown that noise-induced transitions between the states can be minimised by autocorrelating the noise over some critical timescale [3]. The mean extinction time of a metapopulation can be maximised in an analogous fashion by allowing some critical degree of migration between populations [8].

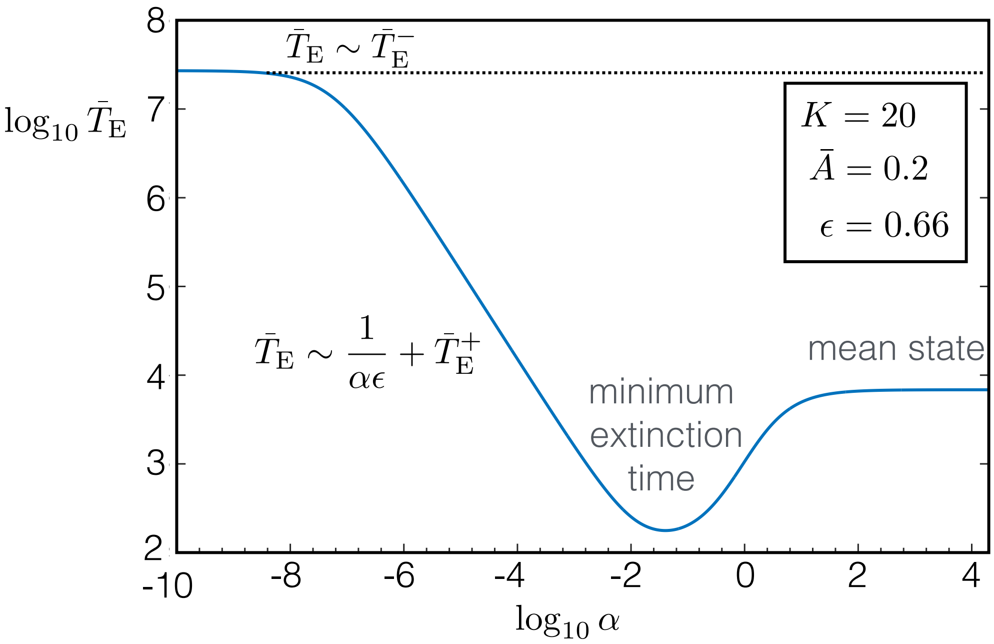

Figure 5: The mean extinction time as a function of the stochastic rate parameter α. Notice the minimum. We explain the origin of the minimum in the text.

The specific cases above are related to a more general problem - the mean first passage time of a particle over a potential barrier whose height fluctuates between a high and a low state [9]. Referred to as "Resonant Activation" (RA), the optimum stochastic frequency was found to be similar to the time it takes to escape when the barrier is at its lowest. Returning to our problem, the analogous conclusion would be that the worst α would correspond with the extinction time in the bad state. We find numerically that this statement is approximately correct at large A (with the agreement being better at larger K). However, for $A > 1$ but not significantly so, a better approximation is to solve for the α that makes the variance equal to the turnover time, or

$$\begin{aligned} \sigma^2 &= \frac{2A\bar{A}a^2}{\alpha(1+\epsilon)^2} \\ &= a \\ \rightarrow \quad \alpha_\mathrm{res} &\approx 2A\bar{A}a\left(1 - \frac{\bar{A}}{A}\right)^{-2}. \end{aligned} \tag{19}$$

We do not yet have a theory to explain this correspondence.

2.5 Origin of "resonance"

In Figure 5, we highlight the dominant influence upon extinction time as a function of α. As $\alpha \rightarrow 0$, the mean extinction time is the mean of the extinction time in the bad state and good state [10]. Owing to the greatly reduced extinction time in the bad compared to the good state, along with $\epsilon < 1$, the mean extinction time in the small α limit simply becomes approximately the extinction time in the good state, $\bar{T}_e^-$.

Increase α such that its inverse $\alpha^{-1} \lesssim \bar{T}_e^-$ and we enter the regime where a population will not typically go extinct before it encounters a bad state. However, once the bad state hits, with high probability, the population goes extinct before another good state is encountered. Accordingly, the mean extinction time in this regime is roughly the mean time to enter a bad state ($1/\epsilon\alpha$), plus the extinction time in the bad state:

$$\bar{T}_e \sim 1/\epsilon\alpha + \bar{T}_e^+. \tag{20}$$

Eventually, α becomes sufficiently large that the most probable trajectory makes it through at least one bad event, recovering during the subsequent good state. Accordingly, the mean extinction time will begin to rise above the extrapolation of $1/\epsilon\alpha + \bar{T}_e^+$. The α at which such a transition occurs corresponds approximately to the case when $\bar{T}_e^+ \approx 1/\alpha$. Substituting this condition into the small-α expression above, we arrive at an estimate for the minimum extinction time

$$\bar{T}_e^{(\min)} \approx \bar{T}_e^+ \left(1 + \frac{1}{\epsilon}\right). \tag{21}$$

Unfortunately, this approximation is only good to within an order of magnitude or so.

One more regime may be described analytically: the limit $\alpha \to \infty$. In this regime, the environment switches so rapidly that the population only "sees" the mean value of environmental forcing $\bar{A}$. Accordingly, the extinction time $\bar{T}_e^{(\infty)}$ corresponds with that which would be calculated in a stationary environment where δ_n is modulated by $\bar{A}n$. Notice that there is no guarantee in general that $\bar{T}_e^{(\infty)} > \bar{T}_e^{(\min)}$, as is apparent from Figure 4. Rather, a minimum exists only if the mean state is sufficiently favourable, or, that increasing α above α_{res} allows the population sufficient time to recover in the good periods, otherwise, increasing α is monotonically more detrimental.

2.6 Implications/Applications

The existence of a minimum extinction time is of potential significance in a variety of ways. First, our results suggest that even stochastic forcing can lead to selection upon a population's reproductive timescale. Recall that all times here are essentially scaled by a, the organism's life cycle turnover rate. If the population is forced at α_{res}, it can increase its mean extinction time by changing a, with an increase in a corresponding with movement to the left in Figure 4. This result is interesting because, naively, one would not expect a population to be able to adapt in this way because its biology cannot 'know' when the next bad event will occur.

On a longer timescale, the existence of a 'most catastrophic catastrophe' at α_{res} implies an important realisation. Of all extinctions throughout evolution history, it appears that to a first approximation, more populations have been removed by fluctuations over roughly the timescale of a life-cycle than any other timescale, though exactly what that timescale is depends upon the species. Considering annual life cycles, fluctuations on a year to year timescale, such as droughts or ice extent, have removed more populations than extremely rare events such as the bollide impact that sealed the dinosaurs' fate. However, as far as recovery goes, the aftermath of a mass extinction probably leads to qualitatively different evolutionary trajectories than occur subsequent to the various background extinctions trickling in as a result of stochastic forcing at α_{res}.

Also of significance, global warming is thought increase the frequency of extreme events. An important issue is knowing how these changes may impact different species. This work suggests that each species' life-cycle timescale is crucial to understanding their extinction risks. For example, as the frequency of events changes, it may enter the resonance of some species, whilst leaving the resonance of others, actually reducing their extinction risks.

The final implication we mention here is with regard to treating infectious diseases. An extinction is good when it means wiping out an illness. Suppose multiple populations are infected by a pathogen, for which treatment is available, but only at a limited supply rate. An important question is how to best distribute the medicine to minimise the extinction time - the time it takes to wipe out the disease. The minimum here may inform how to best go about such treatments.

3 Continuum Limit and Fokker-Planck Approximation

In the form thus far adopted, i.e., discrete population numbers, we are unable to write down any simple, closed-form expressions for extinction times. However, we highlighted several regimes where the extinction time of the full system was dominated by the plus state, the minus state, or some well-defined combination thereof. Analytic approximations for these extinction times may be obtained, but only in large-K limit. In such a limit, we may approximate the population size as a continuum, thereby re-casting the master equation in terms of a Fokker-Planck equation, from which, the mean extinction time may be drawn. It must be cautioned that the continuum results will not be quantitatively the same as the discrete problem [10], but in most cases the general qualitative nature is preserved.

Above, the population may be thought of as occupying one of a semi-infinite number of "steps" n at any one time. Now suppose we define a new variable

$$x \equiv \frac{n}{K}, \tag{22}$$

such that the space between steps is reduced by a factor K. If we make K large, meaning that the equilibrium population size is large, the space of x becomes closer to a continuum. The continuum version of the master equation is derived by considering that probability may flow into position x from other x' at a rate $W(x-x')R(x')\rho(x')$, but flow out at a rate $R(x)\rho(x)$, with $\rho(x)$ taking the place of P_n as the probability density function. It evolves according to the equation

$$\frac{d\rho(x)}{dt} = -R(x)\rho(x) + \int_0^\infty R(x')W(x-x')P(x')dx'. \tag{23}$$

Upon comparison with the birth-death process above, we see that the function $W(x-x')$ forces probability to only flow between nearest neighbours, corresponding to an interval $\Delta x = 1/K$. Accordingly, W takes the form of two Dirac delta functions, one at $x+1/K$ and one at $x-1/K$. We may now integrate equation 23, then perform a Taylor expansion to order $\mathcal{O}(1/K)^2$ about x such that

$$\frac{d\rho(x)}{dt} \approx \frac{\partial}{\partial x}\Bigg[-f(x) + \frac{1}{2K}\frac{\partial}{\partial x}\Big(g^2(x)\Big)\Bigg]\rho(x) \tag{24}$$

where we have defined the "drift" $f(x)$ and "diffusion" $g(x)$ using the continuous analogues of δ_n and β_n above:

$$\begin{aligned} f(x) &= \beta(x) - \delta(x) & g(x)^2 &= \beta(x) + \delta(x) \\ &= ax(1 - \mathcal{A} - x) & &= ax(1 + \mathcal{A} + x). \end{aligned} \tag{25}$$

The Fokker-Planck equation and functions $f(x)$ and $g(x)$ make be applied to the good state by setting $\mathcal{A} = 0$, the bad state using $\mathcal{A} = A$ and the mean, large-α limit using $\mathcal{A} = \bar{A}$. Within each of these regimes, we may calculate an approximate expression for the MTE. However, the functional form of the solution depends upon whether $\mathcal{A} > 1$ or $\mathcal{A} < 1$. We do not consider the case where $\mathcal{A} \sim 1$.

3.1 Asymptotic solutions

In what follows, we solve for the MTE by supposing that the PDF is stationary, but that there exists a current at $x \to \infty$ that is balanced by extinctions at the origin. The equation to to solve is

$$\frac{\partial \rho}{\partial t} = -\frac{\partial}{\partial x}\Bigg[a\,x\,(1 - A - x)\,\rho\Bigg] + \frac{1}{2\,K}\frac{\partial^2}{\partial x^2}\Bigg[a\,x\,(1 + A + x)\,\rho\Bigg] \tag{26}$$

which we rewrite as

$$\begin{aligned} \frac{\partial \rho}{\partial t} &= -\frac{\partial}{\partial x}\Bigg[u(x)\,\rho\Bigg] + \frac{1}{2\,K}\frac{\partial^2}{\partial x^2}\Bigg[v(x)\,\rho\Bigg] \\ &= -\frac{\partial J}{\partial x} \end{aligned} \tag{27}$$

We now solve the steady-state equation to extract the conserved current

$$J = -\frac{1}{2K}\frac{\partial}{\partial x}\Bigg[v(x)\rho(x)\Bigg] + u(x)\rho(x), \tag{28}$$

which has the general solution

$$\rho = -\frac{2KJ}{a}\frac{1}{x(x + 1 + \mathcal{A})}e^{2K\mathcal{F}(x)}\int_0^x e^{-2K\mathcal{F}(x')}dx', \tag{29}$$

where

$$\begin{aligned} \mathcal{F}(x) \equiv \int \frac{u(x)}{v(x)}dx &= \int \frac{1 - \mathcal{A} - x}{1 + \mathcal{A} + x}dx \\ &= \int \Bigg[-1 + \frac{2}{1 + \mathcal{A} + x}\Bigg]dx \\ &= -x + 2\ln(1 + \mathcal{A} + x) \end{aligned} \tag{30}$$

and so the solution for ρ reads

$$\rho = \frac{2KJ}{a} e^{2K(2\ln(1+\mathcal{A}+x)-x)} \frac{1}{x(x+1+\mathcal{A})} \int_0^x e^{-2K(2\ln(1+\mathcal{A}+x')-x')} dx'. \tag{31}$$

The current J is obtained by requiring that the integral over all space of the PDF is unity:

$$\begin{aligned}\int_0^\infty \rho\, dx &= -\frac{2KJ}{a} \int_0^\infty e^{2K(2\ln(1+A+x)-x)} \frac{1}{x(x+1+A)} \int_0^x e^{-2K(2\ln(1+A+x')-x')} dx' dx \\ &= 1\end{aligned} \tag{32}$$

We cannot solve this equation exactly, however, we may make progress by analysing the large K case. Specifically, where $\mathcal{F}(x)$ is positive, $\exp(2K\mathcal{F}(x))$ is very large and vice versa, such that we can approximate which regions of the integral make the largest contribution. The expansions must be carried out separately for $\mathcal{A} > 1$ and $\mathcal{A} < 1$. The reason is that when $\mathcal{A}$ is smaller, ρ will be Gaussian-like about $x = 1 - \mathcal{A}$, where the deterministic dynamics would reach a stead state. However, if $\mathcal{A} > 1$, this argument breaks down because the deterministic dynamics do not possess a steady solution and the PDF becomes pressed up against the origin. The integral above behaves very differently within these two regimes

3.1.1 Calculation for $\mathcal{A} < 1$

We begin with $A < 1$, such that the PDF has a Gaussian-like peak around $x = 1 - A$. The second integrand (the one over x' in 31) is approximately constant within the region of interest (under the Gaussian). More precisely,

$$\begin{aligned}\int_0^x e^{-2K(2\ln(1+\mathcal{A}+x')-x')} dx' &\approx \int_0^x e^{-4K\ln(1+\mathcal{A})} \int_0^\infty e^{-2K\frac{1-\mathcal{A}}{1+\mathcal{A}}x'} dx' \\ &= e^{-4K\ln(1+\mathcal{A})} \frac{1}{2K} \frac{1-\mathcal{A}}{1+\mathcal{A}}.\end{aligned} \tag{33}$$

Therefore, the integral becomes

$$\begin{aligned}1 &\approx -\frac{1}{2K} \frac{1+\mathcal{A}}{1-\mathcal{A}} e^{-4K\ln(1+\mathcal{A})} \frac{2KJ}{a} \int_{-\infty}^\infty \frac{1}{2(1-\mathcal{A})} e^{2K(2\ln(2)-1+\mathcal{A})} e^{-\frac{1}{2}Ky^2} dy \\ &= -\frac{J}{2a} \frac{1+\mathcal{A}}{(1-\mathcal{A})^2} e^{2K(2\ln 2-1+\mathcal{A}-2\ln(1+\mathcal{A}))} \sqrt{\frac{2\pi}{K}} \\ \rightarrow \quad J &= -\sqrt{K} a \frac{(1-\mathcal{A})^2}{1+\mathcal{A}} \sqrt{\frac{2}{\pi}} e^{-2Kc_1} \\ c_1 &\equiv 2\ln 2 - 1 + \mathcal{A} - 2\ln(1+\mathcal{A}),\end{aligned} \tag{34}$$

and so the mean extinction time

$$\boxed{-J^{-1} \approx \frac{1+\mathcal{A}}{a(1-\mathcal{A})^2} \sqrt{\frac{\pi}{2K}} \exp\left(2K[2\ln 2 - 1 + \mathcal{A} - 2\ln(1+\mathcal{A})]\right)} \tag{35}$$

with the most crucial result being the square root combined with an exponential. This expression, for $\mathcal{A} < 1$, is most applicable to the good state and the mean states.

3.1.2 Calculation for $\mathcal{A} > 1$

We return to the original equation to solve:

$$\int_0^\infty \rho dx = -\frac{2KJ}{a}\int_0^\infty e^{2K(2\ln(1+A+x)-x)}\frac{1}{x(x+1+A)}\int_0^{x'} e^{-2K(2\ln(1+A+x')-x')}dx'dx, \tag{36}$$

where, this time, $A > 1$ and so the PDF is pressed up against $x = 0$ (see numerics for P^+). Accordingly, the second integral will receive most of its contribution from $x \approx 0$. The first integral must again be expanded about small x' but **is no longer approximately constant** within the region of interest. Accordingly, we expand the exponent but do not approximate it as linear in x:

$$\begin{aligned}\int_0^x e^{-2K(2\ln(1+A+x')-x')}dx' &\approx \int_0^x e^{-4K\ln(1+A)}e^{-2K\frac{1-A}{1+A}x'}dx' \\ &= e^{-4K\ln(1+A)}\frac{1+A}{1-A}\frac{1}{2K}\left(1-e^{-2K\frac{1-A}{1+A}x}\right).\end{aligned} \tag{37}$$

We deviate once again from the $\mathcal{A} < 1$ case by expanding the exponential term in small x rather than about the maximum in x (because the maximum is now at $x < 0$). Specifically, we approximate

$$e^{2K(2\ln(1+\mathcal{A}+x)-x)} \approx e^{4K\ln(1+\mathcal{A})}e^{2K\frac{1-\mathcal{A}}{1+\mathcal{A}}x} \tag{38}$$

which after substitution yields

$$1 \approx -\frac{J}{a}\frac{1+\mathcal{A}}{1-\mathcal{A}}\int_0^\infty \frac{1}{x(1+\mathcal{A}+x)}\left(e^{2K\frac{1-\mathcal{A}}{1+\mathcal{A}}x}-1\right)dx, \tag{39}$$

which is a relatively simple form as some exponential terms have cancelled out.

Next, we define the positive quantities

$$\begin{aligned}B &\equiv \frac{\mathcal{A}-1}{\mathcal{A}+1} \\ \lambda &\equiv 2K(\mathcal{A}-1)\end{aligned} \tag{40}$$

and a positive variable

$$\xi \equiv 2KBx \tag{41}$$

such that the integral becomes

$$1 \approx \frac{J}{a}2K\int_0^\infty \frac{1}{\xi(\xi+\lambda)}\left(e^{-\xi}-1\right)d\xi. \tag{42}$$

Integrating by parts:

$$\frac{a}{2KJ} \approx \frac{1}{\lambda}{}_0^{\infty}\left[\ln\left(\frac{\xi}{\xi+\lambda}\right)(e^{-\xi}-1)\right] + \frac{1}{\lambda}\int_0^\infty \ln\left(\frac{\xi}{\xi+\lambda}\right)e^{-\xi}d\xi. \tag{43}$$

whence the boundary term vanishes, leading to the compact integral

$$\frac{a}{2KJ} \approx \frac{1}{\lambda}\int_0^\infty \ln\left(\frac{\xi}{\xi+\lambda}\right)e^{-\xi}d\xi. \tag{44}$$

Adding and subtracting $\ln(\lambda)$ yields

$$\begin{aligned}\frac{a\lambda}{2KJ} &= \int_0^\infty \ln\left(\frac{\xi}{\lambda}\right)e^{-\xi}d\xi - \int_0^\infty \ln\left(1+\frac{\xi}{\lambda}\right)e^{-\xi}d\xi \\ &= -\gamma - \int_0^\infty \ln(\lambda)e^{-\xi}d\xi - \int_0^\infty \ln\left(1+\frac{\xi}{\lambda}\right)e^{-\xi}d\xi \\ &\approx -\gamma - \ln\lambda - \int_0^\infty \left(\frac{\xi}{\lambda}+\frac{\xi^2}{2\lambda^2}\right)e^{-\xi}d\xi \\ &= -\gamma - \ln\lambda - \frac{1}{\lambda} - \frac{1}{\lambda^2}. \end{aligned} \tag{45}$$

Finally, the mean extinction time, neglecting $1/K$ terms, becomes

$$\boxed{-J^{-1} \approx \frac{1}{a(\mathcal{A}-1)}\Big(\gamma + \ln(2K(\mathcal{A}-1))\Big)} \tag{46}$$

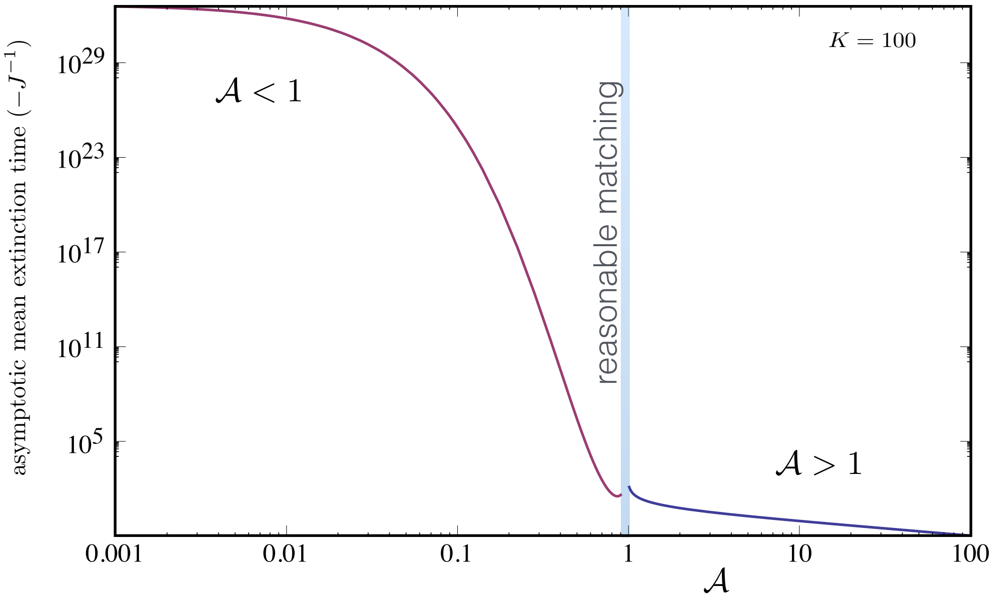

Figure 6: Asymptotic solutions for extinction time in both cases $\mathcal{A} < 1$ and $\mathcal{A} > 1$. We exclude the region near $\mathcal{A}$ as our expansions were not valid there. There is a reasonable degree of matching between the two regimes, which could be improved by utilised an expansion of $\mathcal{A}$ about unity.

Using the above expressions from extinction time, we may approximate the dynamics of the full, environmentally-forced system in the $\alpha \to 0$, small α and $\alpha \to \infty$ limits. Determining whether a minimum extinction time exists or not requires a comparison between the extinction time at α_{res} and the extinction time at $\alpha \to \infty$, which requires an expression for the extinction time near $\mathcal{A} = 1$, as this is where the numerics suggest the minimum disappears. We do not yet have an expression in this regime. However, such a calculation is fairly straight forward and becomes applicable in the case where low extinction times characterise even the mean state.

4 Stochastic Calculus Form

The goal of this section is to obtain a stochastic differential equation that approximates the evolution of the population size, which now takes the form of a stochastic variable X_t[2]. A well-known result is that the probability density function satisfying a Fokker-Planck equation of the form 24, describes a trajectory given by the so-called Itô Stochastic Differential equation

$$dX_t = f(X_t)dt + \left[g(X_t)/\sqrt{K}\,\right] \cdot dW_t, \tag{47}$$

where W_t describes a Weiner Process, whose time derivative produces delta-correlated Gaussian White Noise. Notice that if we suppose $K \to \infty$, the equation aproaches a deterministic differential equation with the population increasing simply as the different between births and deaths ($f(X_t)$).

The form above suggests that the noise associated with random births and deaths is intrinsically white, or at least can be modelled as such. More specifically, there exist two different interpretations of stochastically-forced systems - Itô and Stratonovich. In the Itô interpretation, one considers the system to be forced by noise that exactly satisfies both the Martingale and Markovian conditions [11]. The former implies that the expectation value at some future time equals the current state, whereas the latter suggests that the future state depends only upon the current conditions. The Stratonovich interpretation supposes that no real noise precisely satisfies these criteria and so deriving the noise term requires considering some real noise before taking the white-noise limit. These two interpretations yield different functional forms for the SDE.

Now suppose we wanted to add environmental forcing to the Itô SDE above. Owing to the 'realness' of environmental forcing, it has previously been suggested that one adds a noise term using Stratonovich calculus [12] thereby mixing the interpretations. In what follows, we do not explicitly make this assumption, but rather, we derive an SDE describing the environmentally forced situation, taking particular limits upon α. What we find is an expression that appears to mix the Itô and Stratonovich interpretations, but only within the appropriate limits, suggesting that merely postulating an extra, Stratonovich term is not adequate in most scenarios.

4.1 Mixed interpretations

In order to proceed we must obtain an approximate solution for the full probability

$$P_n = P_n^+ + P_n^-, \tag{48}$$

---

[2] As a convention, the subscript 't' in the field of stochastic calculus refers to 'as a function of time', rather than 'time derivative'.

which is easier at first for the discrete case. To do so, we add the Master equations (10) to obtain

$$\begin{aligned}\frac{dP_n^+}{dt} &= \big(M_{mn} + aAO_{mn} - \alpha I_{mn}\big)P_n^+ + \epsilon\alpha I_{mn}(P_n - P_n^+)\\ \frac{dP_n}{dt} &= M_{nm}P_n + aAO_{nm}P_n^+.\end{aligned} \tag{49}$$

First, we define a matrix

$$C_{mn} \equiv -(M_{mn} + aAO_{mn} - \alpha(1+\epsilon)I_{mn}), \tag{50}$$

such that the master equation for the bad state becomes

$$\frac{dP_n^+}{dt} + C_{mn}P_n^+ = \epsilon\,\alpha\,I_{mn}P_n. \tag{51}$$

For convenience, we now make somewhat of an abuse of notation, using C in favour of C_{mn}, in terms of which, the general solution to the equation above is

$$P_n^+(t) - e^{-C(t-t_0)}P_n^+(t_0) = \epsilon\,\alpha \int_{t_0}^{t} e^{-C(t-s)}P_n(s)ds, \tag{52}$$

where it must be remembered that C is subject to the rules of matrix algebra.

As noted above in the eigenvalue problem, we are interested in the quasi-steady behaviour. Accordingly, it is appropriate to take as our lower bound $t_0 \to \infty$, thereby eliminating the boundary term on the LHS. Furthermore, we may solve the integral on the RHS by way of the following arguments. The matrix C possess a full spectrum of positive eigenvalues, however, by taking the limit $t_0 \to -\infty$, we are essentially stating that the time-evolution of P_n is dominated by the lowest (in magnitude) eigenvalues. Consequently, we may suppose that the integral over all past times s on the RHS obtains the majority of its contribution from the recent past, i.e., small $t-s$. Accordingly, we Taylor expand about $s = t$:

$$P(s) \approx P(t) + (s-t)\frac{dP}{ds}\bigg|_t + \mathcal{O}(s-t)^2. \tag{53}$$

Upon substitution into the general solution (52) we integrate to arrive at the result

$$P_n^+ \approx \epsilon\alpha\bigg(C^{-1}P_n + C^{-2}\frac{dP_n}{dt}\bigg). \tag{54}$$

Finally, we may insert this expression into the master equation for P_n to obtain one single equation for the dynamics, which after some rearranging, takes the form

$$\frac{dP_n}{dt} = \bigg(1 - aA\epsilon\alpha O_{mn}C_{mn}^{-2}\bigg)^{-1}\bigg(M_{mn} + aA\epsilon\alpha O C_{mn}^{-1}\bigg)P_n. \tag{55}$$

We now carry out similar Taylor expansions as before to obtain the continuum approximation to the above equation. However, in addition, we suppose that α and A are large, a condition satisfied by the white-noise limit of our chosen environmental forcing. The final expression, up to terms with second derivatives, is

$$\frac{\partial\rho(x,t)}{\partial t} \approx \frac{\partial}{\partial x}\left[-ax(1-\bar{A}-x)+\frac{1}{2K}\frac{\partial}{\partial x}\left(ax(1+\bar{A}+x)\right)\right]\rho(x,t)+\frac{1}{2}\frac{\partial}{\partial x}\left[\sigma\frac{\partial}{\partial x}\sigma\right]\rho(x,t), \tag{56}$$

where the notation is such that all terms to the left of $\rho(x,t)$ operate upon it. Here, $\sigma^2 = a^2A^2\epsilon/(\alpha(1+\epsilon)^2)$ is the variance of the noise.

There are two crucial things to note about the form above. First, if one truly considers $\alpha = \infty$, then the equation becomes that of the mean state. Importantly, the mean environmental state *acts upon the demographic stochasticity.* However, with α large but finite, an extra term arises on the right that has the form one would infer from a Stratonovich interpretation of white noise with variance σ^2. Accordingly, by prescribing a real source of noise, we could extract the Stratonovich form, but only by making approximations that exclude some of the more interesting aspects of the dynamics, i.e., the the minimum at α_{res}.

Acknowledgements

I would like to thank my two advisors, Charlie Doering and Glenn Flierl for their invaluable mentorship during the Woods Hole GFD Summer School. Charlie's ability to convey detailed concepts was rivalled only by his dedication as a first baseman. Glenn's rational insight always helped simplify problems, even those we hadn't come up with yet. Their combination provided me a very privileged research environment for which I am immensely grateful. Also worth a mention is John Wettlaufer. Throughout the summer, he showed great restraint in keeping a particularly eminent group of professors in line as they vied for attention during seminars. Last but not least I am thankful to Andy Thompson, not just for informing me of the summer school and aiding my application, but also for helping to inspire my interest in geophysical fluid dynamics through two years of excellent teaching and exciting discussion.

References

[1] Shahrezaei, V., Ollivier, J. F., & Swain, P. S. (2008). Colored extrinsic fluctuations and stochastic gene expression. Molecular systems biology, 4(1), 196.

[2] Kamenev, A., Meerson, B., & Shklovskii, B. (2008). How colored environmental noise affects population extinction. Physical review letters, 101(26), 268103.

[3] Assaf, M., Roberts, E., Luthey-Schulten, Z., & Goldenfeld, N. (2013). Extrinsic noise driven phenotype switching in a self-regulating gene. Physical review letters, 111(5), 058102.

[4] Gardiner, C. W. (1985). Handbook of stochastic methods (Vol. 4). Berlin: Springer.

[5] Cairns, B. J., Ross, J. V., & Taimre, T. (2007). A comparison of models for predicting population persistence. ecological modelling, 201(1), 19-26.

[6] Ribalet, F., Swalwell, J., Clayton, S., Jimnez, V., Sudek, S., Lin, Y., ... & Armbrust, E. V. (2015). Light-driven synchrony of Prochlorococcus growth and mortality in the subtropical Pacific gyre. Proceedings of the National Academy of Sciences, 201424279.

[7] Chang, H. H., Oh, P. Y., Ingber, D. E., & Huang, S. (2006). Multistable and multistep dynamics in neutrophil differentiation. BMC cell biology, 7(1), 11.

[8] Khasin, M., Meerson, B., Khain, E., & Sander, L. M. (2012). Minimizing the population extinction risk by migration. Physical review letters, 109(13), 138104.

[9] Doering, C. R., & Gadoua, J. C. (1992). Resonant activation over a fluctuating barrier. Physical review letters, 69(16), 2318.

[10] Doering, C. R., Sargsyan, K. V., & Sander, L. M. (2005). Extinction Times for Birth-Death Processes: Exact Results, Continuum Asymptotics, and the Failure of the Fokker–Planck Approximation. Multiscale Modeling & Simulation, 3(2), 283-299.

[11] Moon, W., & Wettlaufer, J. S. (2014). On the interpretation of Stratonovich calculus. New Journal of Physics, 16(5), 055017.

[12] Hakoyama, H., & Iwasa, Y. (2000). Extinction risk of a density-dependent population estimated from a time series of population size. Journal of theoretical biology, 204(3), 337-359.