On continuum approximations of continuous-time discrete-state stochastic processes of large system size

Davin Lunz

On continuum approximations of continuous-time discrete-state stochastic processes of large system size

Davin Lunz[1,2,3]

[1] *INRIA Saclay – Île de France, 91120 Palaiseau, France*
[2] *École Polytechnique, CMAP 91128 Palaiseau, France*
[3] *Institut Pasteur, 75015 Paris, France*

Abstract

Discrete stochastic processes are an important class of models employed broadly across the sciences. When the system size becomes large, standard approaches can become intractable to solve exactly or even to simulate numerically. It is common to employ continuum approximations that may be more readily solved, and are presumed to converge in the limit as the system size tends to infinity. For example, an expansion of the master equation truncated at second order yields the Fokker–Planck equation, a widely used continuum approximation. Surprisingly, in [5] it is shown that, for birth–death processes, the mean extinction times predicted by the Fokker–Planck approximation may exhibit exponentially large errors, even in the infinite system-size limit. The authors provide "a heuristic argument for the quantitative failure", and conclude that the Fokker–Planck approximation can only be accurate when restrictively tight constraints are placed on the birth and death rates, offering "a warning about the subtlety of the relationship between discrete and continuum approaches". Crucially, however, the underlying source of the inaccuracy has not been addressed. In this paper, we establish a novel quantitative argument characterising how the exponentially large error stems from the finite truncation order. We further argue that this inaccuracy is isolated to a subclass of problems. Thus, the continuum limit is not "delicate" as first thought, nor is the inaccuracy related, in any way, to the birth and death rates. Instead, the continuum limit is robust, the Fokker–Planck approximation is simply a low-order member of a family of superexponentially-converging approximations. Therefore the previously held belief that the continuum approach is only relevant in very special circumstances is to be relaxed: it is uniformly valid and quantifiably accurate. Retaining only a few more terms in the expansion affords very accurate approximations, and should alleviate much of the concern regarding the continuum approach. In establishing the utility of higher-order truncations, this work also contributes to the extensive discussion in the literature regarding the justification of the second-order truncation. While the second-order truncation has appealing features and is sufficiently accurate for many problems, in certain cases it may be necessary to proceed to higher order if additional exponential accuracy is required.

1 Introduction

Discrete stochastic processes in general, and continuous-time Markov chain models in particular, are ubiquitous in physics, chemistry, biology, and throughout the applied sciences [8, 28]. The dynamics of these processes are described by the so-called master equation [9], a first-order system of ordinary differential equations (ODEs) governing the probability of the system to be in any given state at any given time.

In many scenarios, the system size (or population size) is very large. For example, in chemical reactions [22] the number of reactant molecules may be close to Avogadro's constant $\mathcal{O}(10^{23})$, while the number of proteins in a cell [21] can be on the order of $\mathcal{O}(10^6)$, and animal populations can be similarly large. With increasing system size, the master equation becomes prohibitively difficult to solve or simulate numerically, motivating various approximation methods (see [6, 11, 12, 26] and references therein). One common approach is to exploit the large system size to derive a continuum description of the process. This may be achieved by expanding the master equation, and truncating formally negligible terms to obtain a partial differential equation (PDE) describing the evolution of the probability density through space and time; a state-continuum analog of the master equation [9]. Typically, the lowest-order contribution describes the averaged drift, and the next order contribution describes the noise due to

the stochasticity. When higher-order terms are neglected, the resulting equation takes the form of a Fokker–Planck equation [24, Ch. 4].

The expansion truncation at second order finds some justification in the literature. One mathematically appealing reason follows from a result due to Pawula [23], showing that the solution of the PDE obtained when truncating at any finite order beyond the second is not guaranteed to remain non-negative everywhere. Therefore, the quantity may not be interpreted as a probability density. This is not as alarming as it might appear at first sight, since truncation at higher order provides a more accurate approximation. Therefore, even if the solution might become slightly negative, it remains a valid approximation, and may provide better overall predictive power (see [15] and [24, Ch. 4]). Positivity is an appealing property from a modeling perspective, and seems to have motivated much of the literature to ignore higher-order truncations. However, in its absence, the knowledge that the overall error, including negative contributions, are vanishingly small as the truncation order increases is sufficiently reassuring: the convergence promises greater accuracy at the cost of a favourable analytic property.

The second-order truncation is further appealing since it is the minimum truncation order that incorporates the stochastic noise, a key ingredient in many systems [27]. Moreover, the resulting second-order Fokker–Planck PDE belongs to a family of diffusion equations that have been extensively studied [24, Ch. 4], making it familiar and tractable with a catalog of standard and well-known analytical and numerical tools. We offer these as potential reasons, but they are certainly not adequate justification to neglect the consideration of higher-order truncations. Indeed, previous authors [13] have surveyed the literature in search of stronger justification for the second-order truncation, but have not found a convincing way to disentangle the "confused history" of why the second-order truncation is so pervasive.

It is well established [13] that "the relative fluctuations in the time-evolving species populations scale as the inverse square root of the reactant populations". In many contexts this rule is invoked to justify why deterministic equations adequately describe some stochastic process. In light of this, it might be tempting to think that, in the limit of large system size in which we are interested, the influence of stochasticity ought to be neglected entirely. There are two important reasons that explain why such a deterministic description is insufficient for our purposes. Firstly, in section 2 we study mean first-passage time problems, where we determine the time it takes for the process to reach some state nominally separated from a metastable equilibrium state. This excursion from the metastable state eventually occurs (with probability one), even in the presence of arbitrarily weak stochasticity, as small fluctuations accumulate over sufficiently long transients. On the other hand, the deterministic dynamics do not admit any excursions from the metastable state. Mathematically, the large system-size limit and the large time limit do not commute, hence the deterministic systems are insufficient models in the study of these problems. Secondly, in practice, we are interested in being able to accurately describe systems of only moderately large size, where the noise might be asymptotically small but is certainly not to be neglected entirely, as is shown in section 3.

For the sake of concreteness, we now turn our attention to a specific subclass of stochastic processes, known as birth–death processes. However, we emphasise that much of what we say is general, and therefore remains true for a broader class of models, as we will return to. Birth–death processes are a prevalent subclass of stochastic processes [5], describing the continuous-time dynamics of a system of discrete states enumerated by 0, 1, 2, …, Y (where Y can be ∞). Transitions only occur between neighbouring states in the enumeration, so $0 < X < Y$ can only transition to $X - 1$ or $X + 1$. The transition rate from state X to $X + 1$ is called the birth rate, and denoted by $\lambda(X)$, while the transition from X to $X - 1$ is the death rate, and denoted by $\mu(X)$.

In population dynamics models, the birth rate in the state $X = 0$ is typically zero, that is, $\lambda(0) = 0$ and the population has become extinct. The $X = 0$ state is then called an absorbing state. The statistics of the time it takes to reach extinction are of great interest in several such studies [5]. We highlight that, even without this absorbing boundary condition, we may study the passage-time statistics to the $X = 0$ state, or any other state for that matter. As we will explain, our results also pertain to these cases.

The Fokker–Planck approximation has long been considered an accurate model for stochastic processes in the large system-size limit [14]. Perhaps surprisingly, it was shown in [5] that the Fokker–Planck approximation may deviate from the exact discrete solution substantially, producing exponentially large errors in the expected time to extinction, even for arbitrarily large systems. This observation was accompanied by a heuristic argument suggesting that the Fokker–Planck approach is accurate only in the special circumstance that the birth and death rates remain sufficiently close for all states between the equilibrium state and the extinction state. It was thus concluded in [5] that the continuum approach is delicate and only applicable under very strict conditions. It is worth noting that the mean first-passage time deviation of the Fokker–Planck approximation was reported previously in the literature [2, 19].

There are several steps in the derivation of the Fokker–Planck approximation from which error may originate. The continuum description of the population, assuming some regularity of the solutions of the Fokker–Planck equation, are not present in the original discrete setting. Non-analytic terms will not be captured by expanding the master equation [17]. Moreover, the truncation neglects terms, albeit these are formally negligible in the large system-size limit. Previously, the precise source of the exponentially large approximation error was not identified, and therefore not remedied. In this paper, we argue that the continuum approximation, even in enforcing additional structure than present in the discrete description, is exponentially accurate. We present novel calculations showing that the source of the exponentially large errors stems from the finite truncation. The truncated terms characterise the vanishingly small fluctuations, which, even though negligible, can accumulate over long time periods. The calculations demonstrating this error further show that higher-order terms contribute successively less towards the first-passage time, and therefore a truncation is certainly reasonable, but might be required at order greater than two. We also analyse the quasistationary distribution of the same process, and show that it possesses an analogous structure, but due to normalisation of the distribution the influence of the higher-order terms is even less pronounced.

These results contextualise the Fokker–Planck approximation as the lowest-order member of a family of master equation approximations, providing an angle that sheds new light on both the truncation confusion within literature, as well as the Fokker–Planck inaccuracy problem. We deduce that, despite the prevalence of the second-order truncation, higher orders may be desirable in certain contexts, such as the first-passage time problems studied in this paper. Even so, the analysis affirms the robustness of the general continuum approach, negating any concern that the continuum approach is "delicate" or only applicable in "very special situations". In particular, the convergence is not conditional on the birth and death rates, and, while the observations of [5] remain true, we believe that our results should attenuate some of the concern regarding the continuum approach in general. Indeed, the Fokker–Planck approximation deviates from the exact solution, however, this is not due to a failing of the continuum approach, but merely due to it being a low-order truncation. For example, retaining just two extra terms mitigates the deviation significantly.

The rest of the paper is organised as follows. In section 2 we consider the problem of mean first-passage times. In section 2.1, we describe birth–death processes and introduce the mean first-passage problem, outlining the previous results obtained in [5] for a particular birth–death process. In section 2.2 we describe the continuum approximation for arbitrary truncation order N, and demonstrate how a WKB approximation of these truncated systems exhibits vanishingly small error. In section 2.3 we show how the discrete and continuum descriptions agree as a function of the truncation order. In section 3 we proceed to demonstrate how the truncation considerations required in the first-passage time problems pertain to the evolution of the process distribution, by analysing the quasistationary distribution of the same process. In section 3.1 we solve the discrete problem, in section 3.2 we analyse the Fokker–Planck and untruncated approximations, and quantify in section 3.4 how these converge to the discrete solution, highlighting how the higher-order terms are less consequential in this context. Finally, in section 4 we summarise our findings.

2 Mean first-passage time problems

2.1 Discrete birth–death processes

In this section we specify the birth–death processes of interest, and quote results from the literature, which will serve as a benchmark to which we compare the continuum approximations.

The birth–death process is described by the master equation, governing the probability P for the system to be in state X at time t, which is given by

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}P(X,t) = \lambda(X-1)P(X-1,t) - \lambda(X)P(X,t) \\ + \mu(X+1)P(X+1,t) - \mu(X)P(X,t).\end{aligned} \tag{1}$$

Boundary conditions reflect the modeling description that only states $\{X\}_{X=0}^{Y}$ are considered: the death rate for state $X = 0$ must be zero, and, when Y is finite, the birth rate in state $X = Y$ must be zero, that is, $\mu(0) = 0$ and $\lambda(Y) = 0$ if $Y < \infty$.

We are interested in the mean first-passage time to arrive at the state $X = 0$ from any other state,

which we denote by $T(X)$. An equation governing T may be deduced [28, Ch. 12], namely

$$-1 = \lambda(X)T(X+1) - [\lambda(X) + \mu(X)]T(X) + \mu(X)T(X-1), \tag{2}$$

for all processes where the state $X = 0$ is reached with probability 1 from all states. The second-order difference equation (2) is to be closed with two boundary conditions. For the sake of demonstration, we follow [5], and quote the results obtained there without reproducing the calculations. We consider a finite state space $0 \leq X \leq Y$, imposing an absorbing boundary conditions at $X = 0$ and a reflecting boundary condition at $X = Y$, which take the forms

$$T(0) = 0, \qquad T(Y) - T(Y-1) = 1/\mu(Y). \tag{3}$$

The system (2) and (3) admits the solution

$$T(X) = \sum_{i=1}^{X} \left[\frac{1}{\mu(i)} + \left(\prod_{j=1}^{i-1} \frac{\mu(j)}{\lambda(j)} \right) \sum_{k=i+1}^{Y} \frac{1}{\mu(k)} \prod_{\ell=1}^{k-1} \frac{\lambda(j)}{\mu(j)} \right]. \tag{4}$$

Our aim is to compare the solution (4) to the continuum approximation, for large system size, which we denote Ω. While the largest state in the model Y, is not, in principle, the same as the characteristic system size Ω, for the purposes of this paper, we do not distinguish between the two.

It is convenient to introduce scaled rates

$$\lambda(X) = \Omega\bar{\lambda}(X/\Omega), \qquad \mu(X) = \Omega\bar{\mu}(X/\Omega). \tag{5}$$

Following [5] we write $x = X/\Omega$ and consider the concrete case where

$$\bar{\lambda}(x) = \Lambda x(1-x), \qquad \bar{\mu}(x) = x. \tag{6}$$

When $\Lambda > 1$ there is a metastable state $x = x_e := 1 - 1/\Lambda$ where $\bar{\lambda}(x_e) = \bar{\mu}(x_e)$. When $\Lambda < 1$ there are no equilibrium points and the origin is an attractor. The critical value $\Lambda = 1$ is therefore a threshold separating these two qualitatively distinct cases.

As in [5], we focus exclusively on the superthreshold case where $\Lambda > 1$. Using the scalings (5) and the specific forms (6) in the exact solution (4), it is shown in [5] that, in the limit as $\Omega \to \infty$, the solution has a leading order asymptotic form $T(X) \sim T_e$ for $x = X/\Omega = \mathcal{O}(1)$ where

$$T_e = C\mathrm{e}^{\Omega\Phi(x_e)}, \tag{7}$$

and

$$C = \frac{1}{\bar{\lambda}(0) - \bar{\mu}(0)} \sqrt{\frac{2\pi\bar{\lambda}(0)\bar{\mu}(0)}{\Omega[\bar{\lambda}(x_e)\bar{\mu}'(x_e) - \bar{\lambda}'(x_e)\bar{\mu}(x_e)]}} + \mathcal{O}\left(\frac{1}{\Omega}\right), \tag{8}$$

$$\Phi(x) = \int_0^x \log\frac{\bar{\lambda}(\xi)}{\bar{\mu}(\xi)}\,\mathrm{d}\xi. \tag{9}$$

The exponent $\Omega\Phi(x_e)$ in (7) is positive and asymptotically large, therefore it is the central quantity of interest, and any deviation from this value in the continuum approach constitutes an exponentially large error. As such, our primary focus will be on $\log(T)/\Omega$ and, for the purposes of comparison with the continuum approximation, we note from (7) and (8) that

$$\log(T)/\Omega \sim \int_0^{x_e} \log\frac{\bar{\lambda}(\xi)}{\bar{\mu}(\xi)}\,\mathrm{d}\xi. \tag{10}$$

2.2 A family of continuum approximations

In an attempt to reduce system complexity, the continuum approach seeks to approximate the master equation (1) when the characteristic system size is large $\Omega \gg 1$. We adopt the normalised state $x = X/\Omega$ and the scalings (5), and introduce the scaled probability density p

$$p(x,t) = \Omega P(X,t). \tag{11}$$

Substituting p from (11) into the master equation (1), we consider $p(\cdot, t)$ to be defined on the continuum, which justifies a Taylor expansion of the master equation. Truncating the expansion at (possibly infinite) order N, we obtain

$$\frac{\partial}{\partial t}p_N(x,t) = \sum_{k=1}^{N} \frac{(-1)^k}{\Omega^{k-1}k!}\frac{\partial^k}{\partial x^k}\left[\bar{\lambda}(x)p_N(x,t)\right] + \sum_{k=1}^{N} \frac{1}{\Omega^{k-1}k!}\frac{\partial^k}{\partial x^k}\left[\bar{\mu}(x)p_N(x,t)\right]. \tag{12}$$

The Fokker–Planck equation is recovered by setting $N = 2$. We comment that there are two expansions employed in the literature to derive Fokker–Planck approximations: the Kramers–Moyal expansion and van Kampen's system size expansion. Our Taylor expansion is equivalent to the Kramers–Moyal expansion, and we refer the reader to [9, 15, 13] for further discussions in this direction.

We proceed to pose the problem of the mean first-passage time to $x = 0$ for densities governed by (12) starting from any state x, which we denote $\tau_N(x)$. By taking the adjoint of the spatial operator on the right-hand side of (12), it may be shown [28, Ch. 12] that τ_N satisfies the ODE

$$-1 = \bar{\lambda}(x)\sum_{k=1}^{N} \frac{1}{\Omega^{k-1}k!}\frac{\mathrm{d}^k}{\mathrm{d}x^k}\tau_N(x) + \bar{\mu}(x)\sum_{k=1}^{N} \frac{(-1)^k}{\Omega^{k-1}k!}\frac{\mathrm{d}^k}{\mathrm{d}x^k}\tau_N(x) \tag{13a}$$

$$= \sum_{k=1}^{N} \frac{a_k(x)}{\Omega^{k-1}k!}\frac{\mathrm{d}^k}{\mathrm{d}x^k}\tau_N(x), \tag{13b}$$

where

$$a_k(x) = \begin{cases} \bar{\lambda}(x) - \bar{\mu}(x), & k \text{ odd}, \\ \bar{\lambda}(x) + \bar{\mu}(x), & k \text{ even}. \end{cases} \tag{14}$$

Many previous studies avoid the continuum approach, citing the inaccuracy of the Fokker–Planck approximation, and apply a WKB approximation [16, Ch. 7.5] directly of the master equation [19, 1, 20, 29, 4]. We remain within the continuum framework, and solve (13) asymptotically, for $\mathrm{d}\tau_N/\mathrm{d}x$, in the large-system limit as $\Omega \to \infty$ by means of a WKB expansion, whereby

$$\frac{\mathrm{d}}{\mathrm{d}x}\tau_N(x) \sim \left(c_N(x) + \cdots\right)\mathrm{e}^{\Omega V_N(x)}. \tag{15}$$

Then, integrating from $x = 0$ and imposing the boundary condition $\tau_N(0) = 0$, we can write

$$\tau_N(x) = \int_0^x \frac{\mathrm{d}}{\mathrm{d}\xi}\tau_N(\xi)\,\mathrm{d}\xi. \tag{16}$$

It is important to observe from (15) that the dominant contribution from $\mathrm{d}^k\tau_N/\mathrm{d}x^k$ is given by

$$\frac{\mathrm{d}^k}{\mathrm{d}x^k}\tau_N(x) \sim c_N(x)\left[\Omega V_N'(x)\right]^{k-1}\mathrm{e}^{\Omega V_N(x)}, \tag{17}$$

where the prime denotes differentiation. Since the magnitude of the kth derivative is $\mathcal{O}\left(\Omega^{k-1}\right)$, we deduce that, despite each successive term being $\mathcal{O}(\Omega)$ smaller in the approximation (12), their contribution to the exponentially large mean first-passage times are all the same order of magnitude, as $\Omega \to \infty$. Therefore, at leading order, the ODE (13) reduces to the form

$$0 = \sum_{k=1}^{N} \frac{a_k(x)}{k!}V_N'(x)^{k-1}. \tag{18}$$

We see from (18) that $V_N'(x)$ is a real root of a polynomial of degree $(N-1)$. For the Fokker–Planck case, where $N = 2$, we recover

$$V_2'(x) = -2\,\frac{\bar{\lambda}(x) - \bar{\mu}(x)}{\bar{\lambda}(x) + \bar{\mu}(x)}, \tag{19}$$

in agreement with [5].

Our aim is to show that the sequence V_N' converges to Φ', and therefore, the continuum approximations (beyond Fokker–Planck) converge to an exponentially accurate estimate of the underlying discrete process.

Considering first the case of $N = \infty$, we find from (18) that

$$0 = (\bar{\lambda}(x) - \bar{\mu}(x)) \sum_{k=1}^{\infty} \frac{1}{(2k-1)!} V_\infty'(x)^{2k-1} + (\bar{\lambda}(x) + \bar{\mu}(x)) \sum_{k=1}^{\infty} \frac{1}{(2k)!} V_\infty'(x)^{2k} \tag{20a}$$

$$= \left[\bar{\lambda}(x) - \bar{\mu}(x)\right] \sinh V_\infty'(x) + \left[\bar{\lambda}(x) + \bar{\mu}(x)\right] (-1 + \cosh V_\infty'(x)), \tag{20b}$$

where we have identified the hyperbolic trigonometric power series, which, neglecting the trivial solution, is solved by

$$V_\infty'(x) = \log \frac{\bar{\mu}(x)}{\bar{\lambda}(x)}. \tag{21}$$

Note from (18) and the specific forms (6), that the even-indexed coefficients $a_{2k}(x) > 0$ for all x, while the odd-indexed coefficients $a_{2k-1}(x) > 0$ for all $x \in [0, x_e)$. Therefore, $V_N'(x) < 0$ for all $x \in [0, x_e)$ in order for the right-hand side of (18) to be zero. Since V_N' is negative for all $x \in [0, x_e)$, it holds that, on this same domain, $V_N(x_e) < V_N(x)$. Since x_e is a metastable equilibrium point, we expect the mean first-passage time to increase exponentially throughout this domain, except perhaps in a vicinity of the equilibrium point. Therefore, we expect $V_N(x) > 0$ except perhaps near $x \approx x_e$. This motivates us to choose the constant of integration such that $V_N(x_e) = 0$, namely

$$V_N(x) = \int_{x_e}^{x} V_N'(\xi)\, \mathrm{d}\xi. \tag{22}$$

Note that this choice is without loss of generality, since any other constant may be absorbed into the coefficient c. Finally, from the previous observation that $V_N' < 0$ on $x \in [0, x_e)$, we may apply Laplace's method [7, Ch. 2] to the integral (16), to deduce that

$$\log(\tau_N)/\Omega \sim V_N(0) = -\int_0^{x_e} V_N'(\xi)\, \mathrm{d}\xi, \tag{23}$$

where τ_N denotes the mean first-passage time governed by the N-truncated approximation. The calculation is detailed in appendix A.

2.3 Comparing discrete and continuum mean first-passage times

In the Fokker–Planck case of $N = 2$, we substitute (19) into (23) to find that

$$\log(\tau_2)/\Omega \sim 2 \int_0^{x_e} \frac{\bar{\lambda}(\xi) - \bar{\mu}(\xi)}{\bar{\lambda}(\xi) + \bar{\mu}(\xi)}\, \mathrm{d}\xi, \tag{24}$$

in agreement with [5].

In the case of $N = \infty$, we find from (21) and (23) that

$$\log(\tau_\infty)/\Omega \sim \int_0^{x_e} \log \frac{\bar{\lambda}(\xi)}{\bar{\mu}(\xi)}\, \mathrm{d}\xi, \tag{25}$$

which coincides with the asymptotics of the exact solution (10).

We now address intermediate values of $2 < N < \infty$. Note that the series expressions in (20) have an infinite radius of convergence, and therefore we expect these to converge to the value $\log[\bar{\mu}(x)/\bar{\lambda}(x)]$ for all x for which the logarithm is finite. Moreover, over any domain of x for which the logarithm is uniformly bounded, we expect the convergence to be uniform with respect to x. In particular, in the case given by (6), we expect uniform convergence over $x \in (0, 1-\delta)$ for any $\delta > 0$.

Naturally, the convergence requires sufficiently large values of N, however, for small values of N the approximates V_N' may not be well-defined over the entire domain. For example, when $N = 3$, we may solve the quadratic (18) to find that

$$V_3'(x) = \frac{-3\left[\bar{\lambda}(x) + \bar{\mu}(x)\right] \pm \sqrt{3}\sqrt{12\bar{\lambda}(x)\bar{\mu}(x) - 5[\bar{\lambda}(x) - \bar{\mu}(x)]^2}}{2[\bar{\lambda}(x) - \bar{\mu}(x)]}. \tag{26}$$

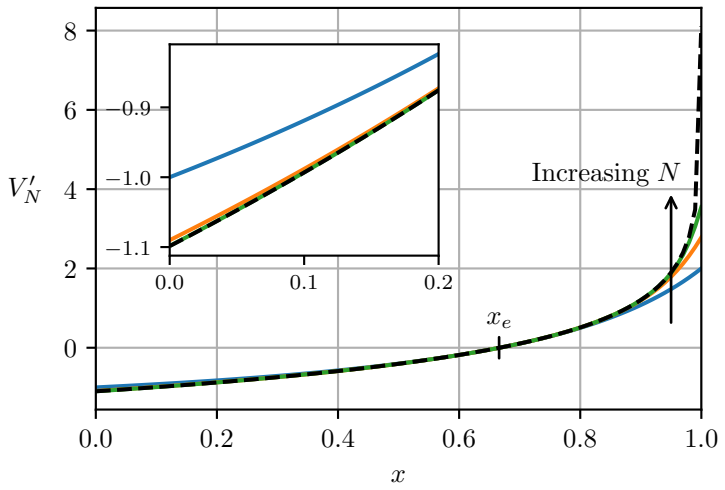

Figure 1: Plots of $V'_N(x)$ for $N \in \{2, 4, 6, \infty\}$. The convergence to the limit of $N = \infty$, plotted as the black dashed curve, is very rapid, near the diverging endpoint $x \approx 1$. The inset shows a zoomed in section of the plot around the endpoint $x \approx 0$.

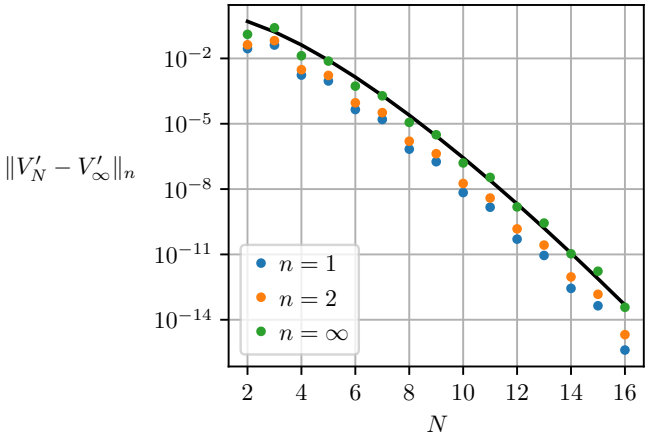

Figure 2: The $L^n([0, 0.9])$ norm of the approximation error of V'_N for increasing N. The black curve shows the predicted convergence behaviour, $1/N!$.

Only one of these roots closely approximates V'_∞ in the domain, however, the more serious problem is that V'_3 is not real for all $x \in [0, 1]$. We can see that, for the forms (6), for $x \approx 1$ we have $\bar{\lambda}(x) \approx 0$, so the first term inside the larger square-root will be arbitrarily small, while the second term is approximately $-5\bar{\mu}(x)^2 \approx -5 < 0$. Our numerical results show that this issue does not arise for even values of N, and for odd values of N, it is relegated to increasingly smaller regions as N increases, as we expect.

To demonstrate this convergence numerically, we use the particular forms in (6), and choose $\Lambda = 3$, since this is the value chosen in [5] to demonstrate the "failure" of Fokker–Planck. We determine $V'_N(x)$ by numerical root-finding of the right-hand side of (18).

First, in fig. 1, we show how the iterates $V'_N(x)$, for even N, rapidly converge to the black dashed limiting curve V'_∞. We see that, as highlighted in [5], the discrepancy of the Fokker–Planck approximation $N = 2$ is noticeable. However, already by $N = 4$ this discrepancy is barely visible over the domain of primary interest $x \in [0, x_e]$. The discrepancy is still visible near the endpoint $x \approx 1$ as V'_∞ diverges, however, this error is of secondary importance since the value of the mean first-passage time is set by V' on $x \in [0, x_e]$, as seen in (23).

Second, we look to quantify the convergence. We expect the truncated series in (20) to have an error of order $\mathcal{O}(1/N!)$ as $N \to \infty$ on any domain where V'_∞ is bounded, say $x \in [0, 1 - \delta]$ for any $\delta > 0$. In fig. 2 we plot the error with $\delta = 0.1$, to find that the convergence matches the predicted rate.

Recall that the odd iterates may not exist everywhere. While all odd iterates do, in fact, exist on the domain $x \in [0, 1-\delta]$ for $\delta = 0.1$ and $\Lambda = 3$, we see that, in particular for small values of N, the error drops much more significantly for even N than for odd N. In fact, the error increases from $N = 2$ to $N = 3$ for this choice of parameters.

In addition to the large system-size limit of $\Omega \to \infty$, we have introduced the new limit of increasing truncation order $N \to \infty$. The asymptotic nature of the limit $\Omega \to \infty$ means that, for large but finite Ω, there might be an optimal truncation point [24, Ch. 4]; some finite N_0 beyond which the error does not decrease for increasing truncation order $N > N_0$.

To recap, we analyse the logarithm of the mean first-passage time for continuum approximations truncated at order N. We find that these approximations converge rapidly to the exact solution with increasing truncation order. While [5] concludes from the discrepancy of the Fokker–Planck approximation of $N = 2$ that the continuum approach is delicate, this analysis offers a different perspective. Even though the low-order truncation of $N = 2$ exhibits some discrepancy in the log mean first-passage time, it certainly captures the qualitative behaviour. Moreover, this discrepancy is significantly reduced for higher truncation order, with the error decreasing superexponentially as $\mathcal{O}(1/N!)$ for large N, with no restrictive conditions needed on the birth or death rates $\bar{\lambda}$ and $\bar{\mu}$. We thus deduce that the continuum approach provides a systematic and robust family of approximations.

We finish by highlighting that these observations do not depend on the particular forms of $\bar{\lambda}$ or $\bar{\mu}$, the restriction that the process be birth–death, or that the process have an absorbing state; these assumptions may all be relinquished. The single crucial ingredient is that the solution of the first-passage problem be of WKB form (15), which leads to higher-order terms in the expansion making $\mathcal{O}(1)$ contributions to the first-passage times. This occurs for processes with a metastable state (or several), and a first-passage to any other state separated from the equilibrium by $\mathcal{O}(\Omega)$ states. Physically, the exponentially long first-passage times are born out of the metastability in combination with small fluctuations: to reach the threshold state requires the rare accumulation of small fluctuations. While these fluctuations are vanishingly small, over a sufficiently large time interval they can (and will) lead to the excursion from the metastable state, and are thus not at all negligible.

Thus, repeating the calculations for a chemical reaction network, for example, with this metastable structure would yield an analogous result. Many popular and useful biological circuits exhibit these properties, for example a genetic toggle switch [10, 25] is a biological system of two mutually repressive genes. The deterministic system [10] exhibits bi-stability, with two stable states representing the dominant production of one gene over the other. Stochasticity is required to drive the switching [25], whereby the continuum approach — if the assumption of moderately large system size is valid — would be an indispensable tool to approximate mean switching times by calculating the mean first-passage time from the metastable states to the invariant manifold separating the regions of attraction of each equilibrium point.

It would be interesting to explore how these observations pertain to higher-order moments of the first-passage time, however, we leave these considerations to future work.

3 Quasistationary distributions

The purpose of this paper is to affirm the usefulness and accuracy of the continuum approach of approximating discrete stochastic processes. Having shown that, for mean first-passage time problems, the exponential error of the Fokker–Planck approximation is recoverable by a higher truncation order, we now analyse the accuracy of the continuum approximations for the problem of distributions governed by the master equation (1). While a similar exponential structure is exposed in this problem, we find that its influence on the distributions is exponentially smaller due to the distributions being normalised. This solidifies the earlier explanation — that the discrepancy observed in [5] is due to the exponentially large passage times, when the accumulation of smaller effects may not be negligible. We find that the Fokker–Planck approximation is remarkably accurate, even for moderate system sizes, as we quantify in this section, and that higher order approximations recover beyond exponential accuracy.

The simplest comparison is between steady state distributions. The unique steady state distribution of the birth–death process using the rates in (6) is extinction. A more interesting distribution is the quasistationary distribution (QSD): that to which the stochastic process settles on intermediate timescales, longer than the birth and death rates characterising the transitions, but shorter than the exponentially long timescales at which extinction is relevant.

3.1 The discrete QSD

We analyse the quasistationary distribution of the discrete stochastic process by not allowing extinction to occur. We set the birth rate at the state $X = 0$ to be positive, that is, $\lambda(0) = \epsilon > 0$, and find that the unique steady distribution of this augmented process is capable of revealing the QSD, by suitably ignoring the $X = 0$ state.

The steady distribution is obtained by setting the left-hand side of the master equation (1) to zero, whereby

$$\begin{aligned}
0 &= -\lambda(0)P(0) + \mu(1)P(1), \\
0 &= \lambda(0)P(0) - (\lambda(1) + \mu(1))P(1) + \mu(2)P(2), \\
&\vdots \\
0 &= \lambda(Y-2)P(Y-2) - (\lambda(Y-1) + \mu(Y-1))P(Y-1) + \mu(Y)P(Y), \\
0 &= \lambda(Y-1)P(Y-1) - \mu(Y)P(Y).
\end{aligned} \tag{27}$$

Successive substitution in (27) shows that $\lambda(X)P(X) = \mu(X+1)P(X+1)$ for all $0 \le X \le Y-1$, from which we deduce that, for all $X > 0$,

$$P(X) = \lambda(0)P(0)\frac{1}{\mu(X)}\prod_{i=1}^{X-1}\frac{\lambda(i)}{\mu(i)}, \tag{28}$$

which determines $P(X)$ for all states $X > 0$. To determine $P(0)$, note that the system of recurrence relations (27) is homogeneous in P, and therefore any solution may be multiplied by a constant to yield another solution. For a valid distribution, we impose $\sum_{X=0}^{Y} P(X) = 1$ which resolves this indeterminacy, and proves that the steady distribution is unique.

In the limit as $\epsilon \to 0$, the steady distribution concentrates all probability mass in the state $X = 0$, that is, extinction, as to be expected. One way of thinking about the QSD is to consider the conditional distribution $P(X \mid X > 0)$. From (28) we see that every state $X > 0$ is multiplied by a factor of $\lambda(0)P(0) = \epsilon P(0)$, but has not other dependence on ϵ or $P(0)$. Therefore, the conditional distribution is simply given by

$$P(X \mid X > 0) = \frac{C}{\mu(X)}\prod_{i=1}^{X-1}\frac{\lambda(i)}{\mu(i)}, \qquad \frac{1}{C} = \sum_{X=1}^{Y}\frac{1}{\mu(X)}\prod_{i=1}^{X-1}\frac{\lambda(i)}{\mu(i)}. \tag{29}$$

The distribution (29) is the QSD. Another way of deriving (29) is to consider the leading-order distribution of the process in the asymptotic limit as $\epsilon \to \infty$.

For the purposes of later comparison, we now deduce an approximate form for the QSD. In [5], it is shown that, in the large system limit $\Omega \gg 1$,

$$\prod_{i=1}^{X-1}\frac{\lambda(i)}{\mu(i)} \sim \frac{\sqrt{\bar{\mu}(0)\bar{\mu}(x)}}{\sqrt{\bar{\lambda}(0)\bar{\lambda}(x)}}\,\mathrm{e}^{\Omega\Phi(x)}, \tag{30}$$

where Φ is defined in (9) and we recall that $x = X/\Omega$. Therefore,

$$P(X \mid X > 0) \sim \frac{\bar{C}}{\sqrt{\bar{\lambda}(x)\bar{\mu}(x)}}\,\mathrm{e}^{\Omega\Phi(x)}, \tag{31}$$

having absorbed all constant terms into the normalising constant $\bar{C}$.

Since $\Omega \gg 1$, the term on the right-hand side of (31), and therefore the QSD, is exponentially dominated by regions where the $\Phi(x)$ takes its maximum value. We note that, with the specific forms (6), the integrand in $\Phi(x)$ takes the form $\log[\bar{\lambda}(x)/\bar{\mu}(x)] = \log[\Lambda(1-x)]$, and is therefore positive on $x \in [0, x_e)$ and negative on $x \in (x_e, 1]$. Therefore $\Phi(x)$ has a unique maximum value at $x = x_e$, and we approximate the QSD via a second-order Taylor expansion, from which we find that the QSD is described by the Gaussian

$$P(X \mid X > 0) \approx c\,\mathrm{e}^{-\Omega\Lambda(x-x_e)^2/2}, \qquad \frac{1}{c} = \int_0^1 \mathrm{e}^{-\Omega\Lambda(\xi-x_e)^2/2}\,\mathrm{d}\xi. \tag{32}$$

Equipped with the exact discrete QSD and an approximate form, we proceed to derive the QSD of the continuum approximation.

3.2 The Fokker–Planck QSD

We begin the continuum approximation of the quasistationary distribution by writing down the Fokker–Planck approximation, obtained by setting $N = 2$ in (12), namely

$$\frac{\partial}{\partial t}p_2(x,t) = -\frac{\partial}{\partial x}\left[\left(\bar{\lambda}(x) - \bar{\mu}(x)\right)p_2(x,t)\right] + \frac{1}{2\Omega}\frac{\partial^2}{\partial x^2}\left[\left(\bar{\lambda}(x) + \bar{\mu}(x)\right)p_2(x,t)\right]. \tag{33}$$

The associated boundary conditions are given by imposing conservation [8, Ch. 5], whereby

$$\left(\bar{\lambda}(x) - \bar{\mu}(x)\right)p_2(x,t) - \frac{1}{2\Omega}\frac{\partial}{\partial x}\left[\left(\bar{\lambda}(x) + \bar{\mu}(x)\right)p_2(x,t)\right] = 0, \tag{34}$$

at the boundary points. With the specific rates (6), the problem (33) and (34) governing p_2 takes the form

$$\frac{\partial}{\partial t}p_2(x,t) = -\frac{\partial}{\partial x}\left[(\Lambda x(1-x) - x)\,p_2(x,t)\right] + \frac{1}{2\Omega}\frac{\partial^2}{\partial x^2}\left[(\Lambda x(1-x) + x)\,p_2(x,t)\right], \tag{35}$$

subject to the boundary conditions

$$(\Lambda x(1-x) - x)\,p_2(x,t) - \frac{1}{2\Omega}\frac{\partial}{\partial x}\left[(\Lambda x(1-x) + x)p_2(x,t)\right] = 0 \text{ at } x = 0, 1. \tag{36}$$

We now examine the Fokker–Planck problem (35) and (36), for which we present exact solutions in the steady case. We also discuss asymptotic approximations, which provide explicit expressions for the transient dynamics, and reveal the underlying structure of the solutions which aids a subsequent convergence analysis.

3.2.1 Exact steady solution

As in the discrete case, an exact solution of (33) and (34), which we denote by p_2, is available in the steady case. Neglecting the time derivative, and integrating (33) once shows that the flux is constant. Then the zero-flux boundary conditions (34) are satisfied at both boundaries, and we are left with the first-order ODE

$$\left(\bar{\lambda}(x) - \bar{\mu}(x)\right)p_2(x) - \frac{1}{2\Omega}\frac{\mathrm{d}}{\mathrm{d}x}\left[\left(\bar{\lambda}(x) + \bar{\mu}(x)\right)p_2(x)\right] = 0. \tag{37}$$

Writing $q = \left(\bar{\lambda} + \bar{\mu}\right)p_2$, we see from (37) that

$$\frac{\mathrm{d}}{\mathrm{d}x}\log(q(x)) = \frac{q'(x)}{q(x)} = 2\Omega\,\frac{\bar{\lambda}(x) - \bar{\mu}(x)}{\bar{\lambda}(x) + \bar{\mu}(x)}, \tag{38}$$

from which it follows that

$$p_2 = \frac{\hat{c}_2\,\mathrm{e}^{\Omega W_2(x)}}{\bar{\lambda}(x) + \bar{\mu}(x)} = \frac{\hat{c}_2\,\mathrm{e}^{\Omega W_2(x)}}{\Lambda x(1-x) + x}, \tag{39}$$

for a normalising constant $\hat{c}_2$, and

$$W_2(x) = 2\int \frac{\bar{\lambda}(x) - \bar{\mu}(x)}{\bar{\lambda}(x) + \bar{\mu}(x)}\,\mathrm{d}x = 2\left(x + \frac{2}{\Lambda}\log[1 + \Lambda(1-x)]\right) + \text{const.} \tag{40}$$

Note that W_2 in (40) is none other than $-V_2$ found in the first-passage time problem (19).

The density (39) asymptotes like $p_2(x) \sim 1/x$ as $x \to 0$, and so is not integrable on $x \in [0, 1]$. This reflects the fact that the steady process accumulates all the probability mass at the absorbing origin. As explored in section 2, we expect this accumulation to be relevant only on exponentially long timescales, and our interest in the quasistationary distribution motivates us to neglect this boundary absorption. We do this by considering the normalised solution on the perturbed domain $[\epsilon, 1]$, for $\epsilon \ll 1$, in analogy with the discrete probability (29) conditioned on being removed from the origin. For $\epsilon \ll 1$ not too small (as we will describe) this accurately describes the quasistationary distribution since it neglects only a small neighbourhood of the origin while guaranteeing the divergent contribution from the neighbourhood of the origin is ignored.

The factor $W_2(x)$ in the exponent in (39) obtains a maximum at the equilibrium point $x = x_e = 1 - 1/\Lambda$. If $\epsilon \gg \mathrm{e}^{-\Omega[W_2(x_e)-W_2(\epsilon)]}$, which is exponentially small since $W_2(x_e) - W_2(0) = \mathcal{O}(1)$ in the limit $\Omega \to \infty$, then $p_2(x_e) \gg p_2(\epsilon)$, whereby we guarantee that the exact solution on the perturbed domain is dominated by the contributions away from the origin. The normalisation constant $\hat{c}_2$ is then given by

$$\frac{1}{\hat{c}_2} = \int_\epsilon^1 \frac{\mathrm{e}^{\Omega W_2(\xi)}}{\Lambda\xi(1-\xi)+\xi}\,\mathrm{d}\xi. \tag{41}$$

The exact solution p_2 in (39) is explicit, up to the normalisation constant $\hat{c}_2$ defined in (41), which may be evaluated by numerical quadrature. Care must be taken to choose the constant of integration for W_2 such that W_2 does not become positive, whereby the calculation becomes susceptible to large errors and numerical overflow, as the integrand extends beyond the range available for standard floating-point calculation. An asymptotic approximation of $\hat{c}_2$ via Laplace's method (detailed in appendix A), produces a closed-form expression that alleviates these problems, namely

$$\begin{aligned}\hat{c}_2 &\sim [\bar{\lambda}(x_e) + \bar{\mu}(x_e)]\,\mathrm{e}^{-\Omega W_2(x_e)}\sqrt{\frac{-\Omega W_2''(x_e)}{2\pi}} \\ &= 2(1-1/\Lambda)\,\mathrm{e}^{-\Omega W_2(x_e)}\sqrt{\frac{\Omega\Lambda}{2\pi}}.\end{aligned} \tag{42}$$

Combining the exponents in (39) and (42) we see that $p_2 \propto \mathrm{e}^{\Omega[W_2(x)-W_2(x_e)]}$, which is no longer exponentially large. In fact, this expression shows that to integrate the exponent from $x = x_e$ is the natural choice of constant, just as discussed in the context of first-passage time problems in the paragraph above (22).

It might be tempting to think that once we have an exact solution there is little more to do. However, the exact solution was only for the steady problem. Moreover, in higher dimensional problems, exact solutions may not be available. For these reasons, we proceed to explore asymptotic solutions. As it turns out, these provide additional quantitative insight even for the steady problem when we have an exact solution.

3.2.2 Asymptotic solution

In this section, our aim is to solve (35) and (36) asymptotically, in the limit as $\Omega \to \infty$. To this end, we expand p_2 asymptotically in inverse powers of Ω

$$p_2 \sim p_{2,0} + \frac{1}{\Omega}p_{2,1} + \cdots. \tag{43}$$

The leading-order form of the governing PDE (35) takes the form

$$\frac{\partial}{\partial t}p_{2,0}(x,t) + \frac{\partial}{\partial x}\left[(\Lambda x(1-x)-x)\,p_{2,0}(x,t)\right] = 0, \tag{44}$$

subject to

$$p_{2,0}(0,t) = p_{2,0}(1,t) = 0. \tag{45}$$

The hyperbolic PDE (44) and (45) may be solved by the method of characteristics. We emphasise that its solution is for the time-dependent problem, and thus the asymptotics provide insight beyond the explicit solutions available only in the steady cases. Characteristic curves satisfy

$$\frac{\mathrm{d}x}{\mathrm{d}t} = \Lambda x(1-x) - x, \tag{46}$$

which admits the general solution

$$x = \left[\frac{1}{x_e} + \left(\frac{1}{x_0} - \frac{1}{x_e}\right)\mathrm{e}^{-(\Lambda-1)(t-t_0)}\right]^{-1}, \tag{47}$$

where the characteristic passes through $x = x_0$ at $t = t_0$.

Along characteristics, the density satisfies

$$\frac{\mathrm{d}}{\mathrm{d}t}p_{2,0}(t) = [2\Lambda x(t) - (\Lambda-1)]\,p_{2,0}(t). \tag{48}$$

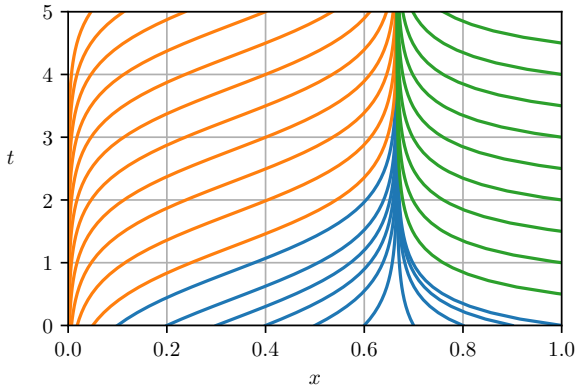

Figure 3: Characteristics (47) for $\Lambda = 3$, emanating from $(x_0, t_0) \in \cup_{i=1}^{3} X_i \times T_i$, where $X_1 = \{j/10\}_{j=1}^{10}$, $X_2 = \{\mathrm{e}^{-j}\}_{j=3}^{13}$, $X_3 = \{1\}$, $T_1 = T_2 = \{0\}$, and $T_3 = \{j/2\}_{j=1}^{9}$. The blue, orange, and green characteristics are associated with $X_i \times T_i$ for $i = 1, 2, 3$, respectively.

Substituting the characteristic form of $x(t)$ from (47) and integrating, we find that, along characteristics, the probability satisfies

$$p_{2,0}(t) = p_{2,0}(t_0)\mathrm{e}^{(\Lambda-1)(t-t_0)}\left[\frac{x_0}{x_e} + \left(1 - \frac{x_0}{x_e}\right)\mathrm{e}^{-(\Lambda-1)(t-t_0)}\right]^{2\Lambda}. \tag{49}$$

For any fixed characteristic with $x_0 > 0$, the term in (49) in square brackets is bounded, and we highlight the fact that the probability density grows exponentially.

We illustrate typical characteristic curves on the (x, t)-plane in fig. 3 for $\Lambda = 3$. We observe that all the characteristics converge towards the equilibrium point $x = x_e = 1 - 1/\Lambda = 2/3$. From (47) we deduce that the convergence is exponential, and from (48) we find that the total density is conserved by exponential growth of the probability density along characteristics. Thus, for any initial condition, the leading-order probability density tends to a delta function, with all the probability mass being concentrated at $x = x_e$.

The leading-order boundary conditions are $p_{2,0} = 0$ at both $x = 0$ and $x = 1$, thus no probability mass enters the domain from the boundaries. This is less important at the left-hand boundary $x = 0$ where the advection vanishes, as illustrated by the orange characteristics that fill the entire upper-left side of the domain, despite all emanating from the x-axis at $t = 0$. However, the green characteristic curves show that the zero boundary condition at $x = 1$ does enter the domain.

The convergence of the leading-order solution towards a delta function does not agree with the discrete QSD (32). The discrepancy is not a matter of proceeding to higher order, since the contribution of higher-order terms to the steady state is beyond all orders. Instead, an inner (boundary) layer of width $\mathcal{O}(1/\sqrt{\Omega})$ develops around $x = x_e$, in which the higher-order terms are not negligible. We may resolve the probability density in the vicinity of this inner layer by adopting the scalings

$$x = x_e + \xi/\sqrt{\Omega}, \qquad p(\xi) = p_{2,0}(x)/\sqrt{\Omega}. \tag{50}$$

Within this inner layer, the equation (35), along with the reinstated second-order terms, takes the form

$$\frac{\partial p}{\partial t} = (\Lambda - 1)\frac{\partial}{\partial \xi}(\xi p) + (1 - 1/\Lambda)\frac{\partial^2 p}{\partial \xi^2}, \tag{51}$$

up to terms of order $\mathcal{O}(1/\sqrt{\Omega})$, and we impose the matching conditions $p \to 0$ as $\xi \to \pm\infty$. Since we only obtained the steady distribution in the discrete case, it suffices to solve the steady inner problem, which admits the solution

$$p = c\,\mathrm{e}^{-\Lambda\xi^2/2}, \tag{52}$$

where c is a normalising constant. Rewriting (52) in terms of the outer probability density by undoing the scalings (50), we find that

$$p_{2,0} = c_{2,0}\,\mathrm{e}^{-\Omega\Lambda(x-x_e)^2/2}, \qquad c_{2,0} \sim \sqrt{\frac{\Omega\Lambda}{2\pi}}, \tag{53}$$

where $c_{2,0} = c\sqrt{\Omega}$ remains a normalising constant, admitting an asymptotic approximation via Laplace's method (see appendix A). Since the outer problem is zero beyond the inner layer, the solution (53) is uniformly valid in x. The Fokker–Planck approximation (53) matches the Taylor-expanded form of the discrete solution (32).

The asymptotic analysis reveals the time-dependent convergence towards the equilibrium, as well as the quasistationary structure around the equilibrium. The asymptotic structure of the outer and inner solutions shows that the deterministic dynamics are dominant beyond the region of width $\mathcal{O}(1/\sqrt{\Omega})$ around the equilibrium, wherein the stochastic dynamics are non-negligible. This scaling will be crucial in estimating the approximation error, not just of the asymptotic solution but also of the exact solution and higher-order solution.

3.3 Higher-order truncations

Both discrete approximations (31) and (32), the exact solution (39), and the asymptotic approximation (53), are all of the form $c(x)\,\mathrm{e}^{\Omega W(x)}$, for various c and W. This suggests that a WKB approximation might be helpful. Indeed, our earlier analysis in section 2 begs the question here: what effect is there on the distribution p_N when truncating the continuum approximation at higher order?

We take the WKB approximation for p_N (not its derivative, as was done in (15)) via

$$p_N \sim (c_N(x) + \cdots)\,\mathrm{e}^{\Omega W_N(x)}. \tag{54}$$

Substituting (54) into the steady version of equation (12), and proceeding just as in section 2.2, one finds that W_N satisfies

$$0 = \sum_{k=1}^{N} \frac{b_k(x)}{k!} W_N'(x)^{k-1}. \tag{55}$$

where $b_k = (-1)^k a_k$ defined in (14). This is analogous to (18), but since the equation governing p_N is the adjoint of that governing τ_N, a sign change in the odd coefficients results.

For Fokker–Planck, where $N = 2$, we recover $W_2 = -V_2$ as above, and proceeding to the next order we find that $c_2(x) = \hat{c}_2 / \left[\bar{\lambda}(x) + \bar{\mu}(x)\right]$. That is, the WKB approximation recovers the exact solution.

In the case of $N = \infty$, and following the algebra from section 2, we find that

$$W_\infty' = \log \frac{\bar{\lambda}(x)}{\bar{\mu}(x)} = -V_\infty' = \Phi', \tag{56}$$

in precise agreement with the discrete approximation (30). Despite it being subdominant, for the sake of completeness we calculate $c_\infty(x)$ by proceeding to the next order (the calculations are performed in appendix B), giving the approximation

$$p_\infty \sim \frac{\hat{c}_\infty}{\sqrt{\bar{\lambda}(x)\bar{\mu}(x)}}\,\mathrm{e}^{\Omega\Phi(x)}, \qquad\qquad \hat{c}_\infty \sim (1 - 1/\Lambda)\,\mathrm{e}^{-\Omega\Phi(x_e)}\sqrt{\frac{\Omega\Lambda}{2\pi}}. \tag{57}$$

We note that in this case, as for the Fokker–Planck case of $N = 2$, the density p_∞ is not integrable at the origin since $p_\infty \sim 1/x$ as $x \to 0$. Despite there also being a singularity at $x = 1$, it is removable since $p_\infty \sim \sqrt{1 - x}$ as $x \to 1$. Again we see from (57) that $x = x_e$ is the natural place to integrate from, however, we preserve the definition of Φ taken from [5].

3.4 Comparing discrete and continuum QSDs

Let us recap the expressions we determined for the QSD. Given the exact discrete solution P, we derived the closed-form approximation (31), and, Taylor expanding the exponent to second order, a further approximation (32). For the continuum family of approximations indexed by N, we obtained an exact Fokker–Planck ($N = 2$) solution (39), denoted p_2, an asymptotic Fokker–Planck solution (53), denoted $p_{2,0}$, and a WKB approximation of the $N = \infty$ case (57), denoted p_∞. In this section, we compare these expressions and draw conclusions by comparing and contrasting these results to the first-passage time results of section 2.

The discrete approximation (31) matches exactly with p_∞ from (57). The Taylor expansion of the former (32) matches precisely with the asymptotic Fokker–Planck solution $p_{2,0}$ from (53). As with the

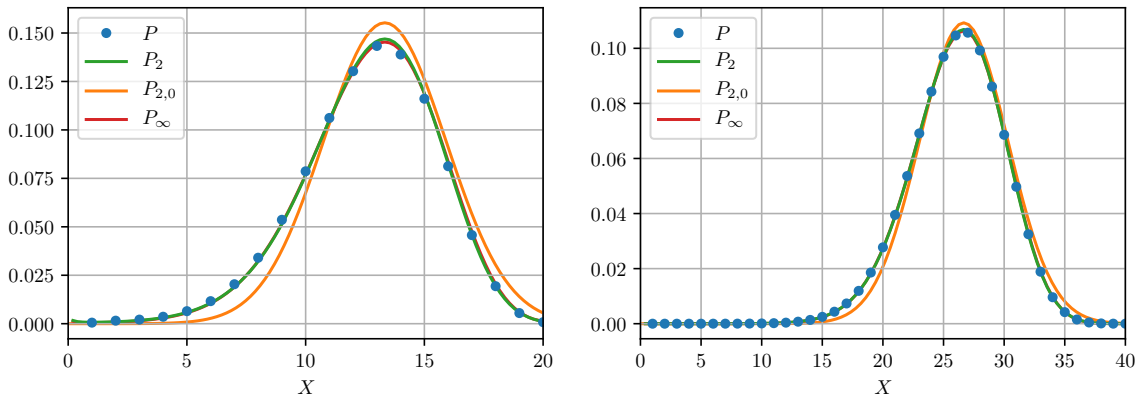

(a) $\Omega = 20$. (b) $\Omega = 40$.

Figure 4: A comparison of the QSDs. Exact discrete solution P given by (29), the exact Fokker–Planck solution P_2 given in (39), the asymptotic Fokker–Planck solution $P_{2,0}$ given in (53), and the WKB approximation of the untruncated equation P_∞ given in (57). The comparison is done in the original scaling, hence the capital letters, by undoing (11).

first-passage time problem, despite being formally smaller, all higher orders contribute to the exponent. In contrast to the first-passage time case, here this discrepancy is particularly negligible. This is because the probabilities are normalised, therefore the probability density is exponentially small beyond a neighbourhood around the equilibrium point $x = x_e$, and local to the equilibrium, the exponents agree to second order. The outcome is that the errors associated with finite truncation are asymptotically small. We proceed to demonstrate this concretely by analysing the error asymptotics and demonstrating convergence, noting the quality of the approximations even for moderate system sizes Ω.

For the sake of measuring the accuracy of the approximations, we compare them to the exact discrete solution P given in (29). To this end, we undo the scaling (11) and denote the probability densities in capital letters. In fig. 4, we illustrate the distribution for different moderately large values of Ω. The solutions P_2 and P_∞ agree remarkably well even for $\Omega = 20$, being almost indistinguishable. We see in fig. 4a good qualitative agreement for the asymptotic approximation $P_{2,0}$, and quantitatively improved agreement in fig. 4b. For smaller values of Ω the quasistationary distribution deviates noticably from the Gaussian distribution. It is worth emphasising that the typical system size — a natural parameter choice to scale the system with — is more accurately given in these cases by $x_e\Omega = (2/3)\Omega$, which makes the agreement even more impressive.

From the matched asymptotic expansion in section 3.2.2 we deduce that $\mathcal{O}(\sqrt{\Omega})$ states in the vicinity of the equilibrium will contribute significantly to the probability density, while states beyond this region will be exponentially small. This allows us to analyse errors local to the equilibrium, which we do in three stages.

First, we compare the matched asymptotic Fokker–Planck solution $P_{2,0}$ to the exact Fokker–Planck solution P_2. The boundary-layer analysis performed to produce (53) is formally accurate up to $\mathcal{O}(1/\Omega)$. We thus predict that the L^n-norm of the discrepancy should asymptote like

$$\begin{aligned} \|P_{2,0} - P_2\|_n = \left(\sum_X \left[P_{2,0}(X) - P_2(X)\right]^n\right)^{1/n} &= \mathcal{O}\left(\left[\sqrt{\Omega}\,\frac{1}{\Omega^n}\right]^{1/n}\right) \\ &= \mathcal{O}\left(\Omega^{-1+1/(2n)}\right). \end{aligned} \tag{58}$$

Second, we compare the exact Fokker–Planck solution P_2 and the WKB approximation of the untruncated system P_∞, taking only contributions from $|x - x_e| = \mathcal{O}(1/\sqrt{\Omega})$ into account. Since their exponents $W_2(x) - W_2(x_e)$ and $\Phi(x) - \Phi(x_e)$ agree locally up to second order, their difference is of the form $[W_2 - W_2(x_e)] - [\Phi - \Phi(x_e)] = [W_2'''(\xi) - \Phi'''(\xi)](x - x_e)^3/6$ for some ξ. We thus deduce that

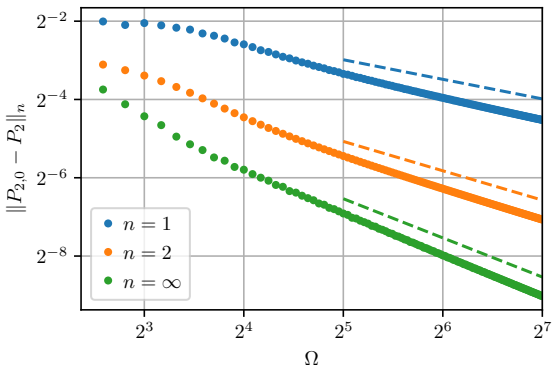

(a) Slope $-1 + 1/(2n)$ from (58).

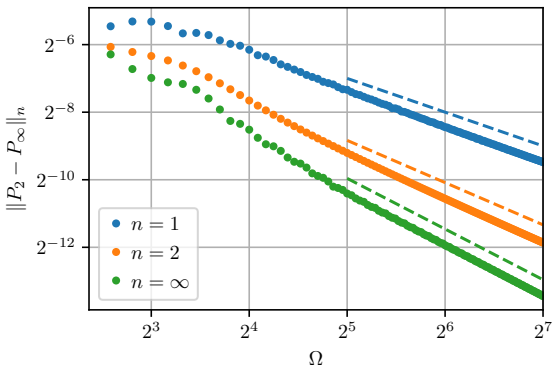

(b) Slope $-1.5 + 1/(2n)$ from (60).

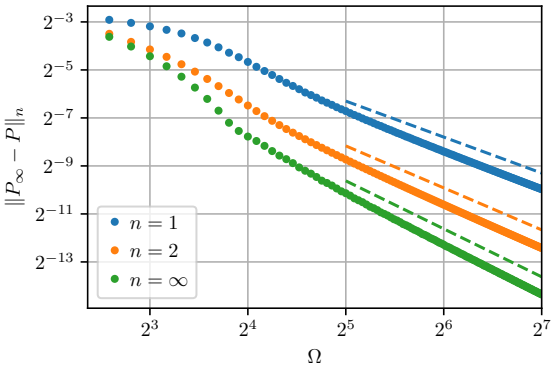

(c) Slope $-2 + 1/(2n)$ from (61).

Figure 5: Discrepancies between different QSDs. Dashed lines represent the predicted convergence. The comparison is done in the original scaling, hence the capital letters, by undoing (11).

$\Omega([W_2(x) - W_2(x_e)] - [\Phi(x) - \Phi(x_e)]) = \mathcal{O}(1/\sqrt{\Omega})$, from which it follows that

$$p_2 - p_\infty \sim \left[\frac{2}{\bar{\lambda}(x) + \bar{\mu}(x)} \mathrm{e}^{\Omega[W_2(x) - W_2(x_e) - \Phi(x) + \Phi(x_e)]} - \frac{1}{\sqrt{\bar{\lambda}(x)\bar{\mu}(x)}}\right] \times x_e \sqrt{\frac{\Omega\Lambda}{2\pi}} \mathrm{e}^{\Omega[\Phi(x) - \Phi(x_e)]} \tag{59a}$$

$$\sim \left[\frac{2}{\bar{\lambda}(x) + \bar{\mu}(x)} - \frac{1}{\sqrt{\bar{\lambda}(x)\bar{\mu}(x)}}\right] x_e \sqrt{\frac{\Omega\Lambda}{2\pi}} \mathrm{e}^{\Omega[\Phi(x) - \Phi(x_e)]} \tag{59b}$$

$$= -\left[\frac{\left(\sqrt{\bar{\lambda}(x)} - \sqrt{\bar{\mu}(x)}\right)^2}{\sqrt{\bar{\lambda}(x)\bar{\mu}(x)}\left[\bar{\lambda}(x) + \bar{\mu}(x)\right]}\right] x_e \sqrt{\frac{\Omega\Lambda}{2\pi}} \mathrm{e}^{\Omega[\Phi(x) - \Phi(x_e)]}. \tag{59c}$$

Fascinatingly, the difference between these two solutions essentially boils down to the difference between (the inverses of) the arithmetic and geometric means of $\bar{\lambda}$ and $\bar{\mu}$, encapsulated in the bracketed term in (59b). Since $\sqrt{\bar{\lambda}(x)} - \sqrt{\bar{\mu}(x)} \sim \Lambda\sqrt{x}(x - x_e)/2 = \mathcal{O}(1/\sqrt{\Omega})$, we conclude that

$$\|P_2 - P_\infty\|_n = \mathcal{O}\left(\Omega^{-1.5+1/(2n)}\right). \tag{60}$$

Finally, we compare P_∞ and the exact discrete solution P. The next order correction in WKB approximation is $O(1/\Omega)$, which suggests that

$$\|P_\infty - P\|_n = \mathcal{O}\left(\Omega^{-2+1/(2n)}\right). \tag{61}$$

In fig. 5 we plot each of the continuum approximations compared with the exact discrete solution. We find good agreement with the error estimates derived above. Since each successive error estimate was

smaller than the previous, we may telescope the errors and deduce the same error in comparison with the discrete solution P. For example, we see that $\|P_{2,0} - P\|_n \leq \|P_{2,0} - P_2\|_n + \|P_2 - P_\infty\|_n + \|P_\infty - P\|_n = \mathcal{O}(\Omega^{-1+1/(2n)})$, since the first term on the right-hand side of the inequality is dominant. Even though the asymptotic Fokker–Planck solution is the least accurate, it was invaluable in uncovering the structure that underpins the error estimates for all of the solutions.

To conclude, we find that the probability distribution contains a similar structure to the mean first-passage times: higher-order terms in the continuum master equation approximation are, formally, increasingly negligible, and yet persist in contributing to the leading-order exponent. However, in contrast to the first-passage time context, a quantitative analysis explains why this discrepancy is even less important for the probability distribution. Thus the Fokker–Planck approximation proves remarkably accurate at resolving the quasistationary distribution of the discrete process, and provides valuable transient and structural information about the problem, while higher-order contributions make only an asymptotically small contribution.

4 Conclusions

Following observations of exponentially large errors in mean first-passage time problems for birth–death processes, it has been suggested that the Fokker–Planck approximation is only valid under strict conditions on the birth and death rates, as well as the system size being large but not too large, resulting in uncertainty surrounding the accuracy and applicability of the technique.

In this paper we present a quantitative argument explaining the cause of the discrepancy. The Fokker–Planck approximation is arrived at via truncation of the continuous approximation of the master equation at the second order. In mean first-passage time problems, each term in this series makes an order unity contribution (with respect to the system size limit) to the first-passage time, therefore low-order truncation should be expected to lead to exponential error. However, the terms in the series diminish rapidly (at a factorial rate, with respect to the truncation order), and so corrections to the Fokker–Planck approximation within the continuum framework are available and tractable. This is the appropriate perspective from which to view the approximation error: Fokker–Planck is the coarsest member in a hierarchy of convergent approximations. Therefore, we expect it to be an invaluable tool, even for first-passage time problems, if only qualitatively. In instances where more quantitative agreement is necessary, the appropriate higher-order truncations are available.

This perspective — viewing Fokker–Planck merely as a member of a family of continuum approximations — helps, in part, to demystify the confusion surrounding why truncation at second order, but not beyond, is so ubiquitous. The reality is, simply, that there is no justification to truncate only at second order. Since [3] "all models are wrong the scientist must be alert to what is importantly wrong", we reiterate the viewpoint that the non-negativity guaranteed by the second order truncation is not necessarily of greater importance than the approximation accuracy. Then, while the Fokker–Planck approximation is sufficiently accurate in many cases, should the need arise, other members of the continuum family may provide increased accuracy.

The exponentially large error of the Fokker–Planck approximation is not as dire as it may appear. To emphasise how these higher-order terms contribute contribute in other problems, we analyse the quasistationary distribution of the same birth–death process. Despite an analogous structure, the Fokker–Planck approximation is exceptionally accurate, even for moderate system sizes, and provides a quantification of where deterministic versus stochastic contributions dominate the dynamics. We thus confirm that the previously observed inaccuracies pertain solely to problems of highly specific structure as revealed through the asymptotic analysis of the exponentially long first-passage times.

We perform the calculations for the specific birth–death process analysed in [5] for the sake of concreteness. However, our method is generic and does not rely on specific features of the birth–death process, or the particular forms of the birth or death rates. The analysis merely requires the existence of a metastable equilibrium state from which the process deviates via small fluctuations, and the first-passage time to a state at distance $\mathcal{O}(\Omega)$ from the metastable state. We emphasise that the first-passage time may be considered to a non-absorbing state; in fact, no absorbing state need be present at all. We thus conclude that our observations regarding the continuum approximation in general, and the Fokker–Planck approximation in particular, generalise to stochastic processes relevant in other fields, such as chemical reaction networks.

The practitioner should be aware of the limits of any tool used, and the continuum approach to approximate discrete stochastic processes is no exception. For further discussion in this direction we refer the reader to [13, 2]. While we agree with the observations that the continuum approach has

some approximation error, we hope that this contribution serves to reaffirm the utility of the continuum family of approximations presented, including Fokker–Planck, by demonstrating their general accuracy and explanatory capability, while also establishing a quantitative explanation and remedy of error in first-passage time problems.

A Laplace's method

In this appendix, we apply Laplace's method to obtain an asymptotic approximation of integrals of the form $\int g(\xi)\,\mathrm{e}^{\Omega f(\xi)}\,\mathrm{d}\xi$ in the limit as $\Omega \to \infty$, assuming that both $g(x)$ and $f(x)$ are of order unity. The central idea is that, since the exponent is large, the dominant contributions to the integral come from neighbourhoods of global maxima of f, as all other regions are exponentially smaller. We demonstrate two cases from the main text: first, where f has a global maximum on the boundary of the domain (with non-zero derivative), and second, where f has a global maximum in the domain interior.

A.1 Boundary maximum

We begin with (15) and (16) for the mean first-passage time with truncation order N, whereby

$$\tau_N = \int_0^x \frac{\mathrm{d}}{\mathrm{d}\xi}\tau_k(\xi)\,\mathrm{d}\xi \sim \int_0^x c(\xi)\,\mathrm{e}^{\Omega V_N(\xi)}\,\mathrm{d}\xi, \tag{62}$$

with $V_N' < 0$ on $x \in [0, x_e)$. Therefore, the dominant contribution is expected to be in a neighbourhood of the origin, and we split the domain of integration via

$$\tau_N \sim \int_0^\delta c(\xi)\,\mathrm{e}^{\Omega V_N(\xi)}\,\mathrm{d}\xi + \int_\delta^x c(\xi)\,\mathrm{e}^{\Omega V_N(x)}\,\mathrm{d}\xi, \tag{63}$$

for some $\delta \ll 1$ which we will determine. Expanding locally allows us to evaluate each integral contribution as follows. Near the origin, and assuming that

$$\Omega\delta \gg 1 \qquad \text{and} \qquad \Omega\delta^2 \ll 1, \tag{64}$$

we find that

$$\begin{aligned}
\int_0^\delta c(\xi)\,\mathrm{e}^{\Omega V_N(\xi)}\,\mathrm{d}\xi &= \int_0^\delta [c(0) + \mathcal{O}(\delta)]\,\mathrm{e}^{\Omega[V_N(0)+\xi V_N'(0)+\mathcal{O}(\delta^2)]}\,\mathrm{d}\xi && \text{(65a)}\\
&\sim c(0)\,\mathrm{e}^{\Omega V_N(0)} \int_0^\delta \mathrm{e}^{\Omega\xi V_N'(0)}\,\mathrm{d}\xi && \text{(65b)}\\
&= -\frac{c(0)}{\Omega V_N'(0)}\,\mathrm{e}^{\Omega V_N(0)}\left[1 - \mathrm{e}^{\Omega\delta V_N'(0)}\right] && \text{(65c)}\\
&\sim -\frac{c(0)}{\Omega V_N'(0)}\,\mathrm{e}^{\Omega V_N(0)}. && \text{(65d)}
\end{aligned}$$

Away from the origin, since $V_N' < 0$, we find that the integral is exponentially smaller than the contribution (65), admitting the bound

$$\int_\delta^x c(\xi)\,\mathrm{e}^{\Omega V_N(x)}\,\mathrm{d}\xi = \mathcal{O}\left(\mathrm{e}^{\Omega V_N(\delta)}\right) = \mathcal{O}\left(\mathrm{e}^{\Omega[V_N(0)+\delta V_N'(0)+\mathcal{O}(\delta^2)]}\right). \tag{66}$$

It thus suffices to find δ satisfying conditions (64), for example, $\delta = \Omega^{-\alpha}$ for $\alpha \in (1/2, 1)$. Then, substituting (65) and (66) into (62), we deduce that

$$\tau_N \sim -\frac{c(0)}{\Omega V_N'(0)}\,\mathrm{e}^{\Omega V_N(0)}, \tag{67}$$

from which we deduce that

$$\log(\tau_N)/\Omega \sim V_N(0). \tag{68}$$

A.2 Interior maximum

In this subsection, we begin with (41), whereby we have the probability distribution p_2 normalised by the constant $\hat{c}_2$, such that

$$\frac{1}{\hat{c}_2} = \int_\epsilon^1 \frac{\mathrm{e}^{\Omega W_2(\xi)}}{\bar{\lambda}(\xi)+\bar{\mu}(\xi)}\,\mathrm{d}\xi\,\mathrm{d}\xi, \tag{69}$$

where $W_2(x)$ obtains a unique maximum at $x = x_e$. We split the domain of integration, isolating a neighbourhood of the maximum $[x_e-\delta, x_e+\delta]$ for $\delta \ll 1$ that we will describe. Assuming

$$\Omega\delta^2 \gg 1 \qquad \text{and} \qquad \Omega\delta^3 \ll 1, \tag{70}$$

we see that

$$\begin{aligned}
\int_{x_e-\delta}^{x_e+\delta} \frac{\mathrm{e}^{\Omega W_2(\xi)}}{\bar{\lambda}(\xi)+\bar{\mu}(\xi)}\,\mathrm{d}\xi\,\mathrm{d}\xi &= \int_{x_e-\delta}^{x_e+\delta} \frac{1+\mathcal{O}(\delta)}{\bar{\lambda}(x_e)+\bar{\mu}(x_e)}\,\mathrm{e}^{\Omega[W_2(x_e)+W_2''(x_e)(\xi-x_e)^2/2+\mathcal{O}(\delta^3)]}\,\mathrm{d}\xi && \text{(71a)}\\
&\sim \frac{\mathrm{e}^{\Omega W_2(x_e)}}{\bar{\lambda}(x_e)+\bar{\mu}(x_e)} \int_{x_e-\delta}^{x_e+\delta} \mathrm{e}^{\Omega W_2''(x_e)(\xi-x_e)^2/2}\,\mathrm{d}\xi && \text{(71b)}\\
&= \frac{\mathrm{e}^{\Omega W_2(x_e)}}{\bar{\lambda}(x_e)+\bar{\mu}(x_e)} \sqrt{\frac{2}{-\Omega W_2''(x_e)}} \int_{-\zeta_0}^{\zeta_0} \mathrm{e}^{-\zeta^2}\,\mathrm{d}\zeta && \text{(71c)}\\
&\sim \frac{\mathrm{e}^{\Omega W_2(x_e)}}{\bar{\lambda}(x_e)+\bar{\mu}(x_e)} \sqrt{\frac{2\pi}{-\Omega W_2''(x_e)}}, && \text{(71d)}
\end{aligned}$$

where $\zeta_0 = \delta\sqrt{-\Omega W_2''(x_e)/2}$. The deduction of (71d) follows by extending the range of integration to $\zeta = \pm\infty$, which introduces exponentially small errors assuming conditions (70) are met.

Away from the maximum, $W_2'(x) \geq 0$ on $x < x_e$, allowing us to bound the integral via

$$\int_\epsilon^{x_e-\delta} \frac{\mathrm{e}^{\Omega W_2(\xi)}}{\bar{\lambda}(\xi)+\bar{\mu}(\xi)}\,\mathrm{d}\xi = \mathcal{O}\left(\mathrm{e}^{\Omega W_2(x_e-\delta)}\right) = \mathcal{O}\left(\mathrm{e}^{\Omega[W_2(x_e)-\delta^2 W_2''(x_e)/2+\mathcal{O}(\delta^3)]}\right), \tag{72}$$

which is exponentially smaller than the contributions near the maximum (71), when conditions (70) are satisfied. Similarly, $W_2'(x) \leq 0$ on $x > x_e$, and the associated integral contribution is exponentially smaller than contribution (71).

Finally, conditions (70) are satisfied by, for example, $\delta = \Omega^{-\alpha}$ for $\alpha \in (1/3, 1/2)$. We therefore conclude that

$$\hat{c}_2 \sim \left[\bar{\lambda}(x_e)+\bar{\mu}(x_e)\right]\mathrm{e}^{-\Omega W_2(x_e)}\sqrt{\frac{-\Omega W_2''(x_e)}{2\pi}}. \tag{73}$$

B First-order calculation for WKB coefficient

The steady form of equation (12) for $N = \infty$ is given by

$$0 = \sum_{k=1}^{\infty} \frac{1}{\Omega^{k-1}k!}\frac{\mathrm{d}^k}{\mathrm{d}x^k}\left[b_k(x)p_\infty(x)\right], \qquad b_k(x) = \begin{cases} \bar{\mu}(x)-\bar{\lambda}(x), & k \text{ odd},\\ \bar{\mu}(x)+\bar{\lambda}(x), & k \text{ even}. \end{cases} \tag{74}$$

Taking the WKB ansatz (and suppressing the subscript $N = \infty$ to simplify the notation)

$$p_\infty \sim (c(x)+o(1))\,\mathrm{e}^{\Omega W(x)}, \tag{75}$$

we seek to express the summand term $\mathrm{d}^k(b_k p)/\mathrm{d}x^k$ in successive powers of Ω. We employ two main results. First, Liebniz's product rule, whereby

$$\frac{\mathrm{d}^k}{\mathrm{d}x^k}[f(x)g(x)] = \sum_{i=0}^{k}\binom{k}{i} f^{(k-i)}(x)g^{(i)}(x). \tag{76}$$

Second, Faà di Bruno's formula [18], namely

$$\frac{\mathrm{d}^k}{\mathrm{d}x^k} f(g(x)) = \sum_{m_1+2m_2+\cdots+km_k=k} \frac{k!}{m_1!\cdots m_k!} f^{\left(\sum_{i=1}^k m_i\right)}(g(x)) \prod_{j=1}^{k} \left(\frac{g^{(j)}(x)}{j!}\right)^{m_j}, \tag{77}$$

where the outer sum is taken over all k-tuples of natural numbers satisfying the constraint $m_1 + 2m_2 + \cdots + km_k = k$. Using $f(x) = \mathrm{e}^{\Omega x}$ and $g(x) = W(x)$, we see that in (77) the order of magnitude of each summand is determined by the order of the derivative of f. The lowest order contribution comes from the k-tuple $(k, 0, \ldots, 0)$, the first order contribution comes from $(k-2, 1, 0\ldots, 0)$, while contributions from all other admissible k-tuples are of higher order. It then follows from (77) that

$$\frac{\mathrm{d}^k}{\mathrm{d}x^k} \mathrm{e}^{\Omega W(x)} = \mathrm{e}^{\Omega W} \left[\Omega^k (W')^k + \Omega^{k-1} \frac{k(k-1)}{2} (W')^{k-2} + \mathcal{O}\left(\Omega^{k-2}\right)\right]. \tag{78}$$

Then from Leibniz's formula (76) and (78) we see that

$$\begin{aligned}
\frac{\mathrm{d}^k}{\mathrm{d}x^k} p_\infty &= \sum_{i=0}^{k} \binom{k}{i} c^{(k-i)} \frac{\mathrm{d}^i}{\mathrm{d}x^i} \mathrm{e}^{\Omega W(x)} \\
&= \mathrm{e}^{\Omega W(x)} \left\{ \Omega^k \left[(W')^k c\right] \right. \\
&\qquad \left. + \Omega^{k-1} \left[\frac{k(k-1)}{2} (W')^{k-2} W'' c + k (W')^{k-1} c'\right] + \mathcal{O}\left(\Omega^{k-2}\right) \right\}.
\end{aligned}$$

An additional application of Leibniz's formula produces

$$\begin{aligned}
\frac{\mathrm{d}^k}{\mathrm{d}x^k} [b_k p_\infty] &= \sum_{i=0}^{k} \binom{k}{i} b_k^{(k-i)} \frac{\mathrm{d}^i}{\mathrm{d}x^i} p_\infty \\
&= \mathrm{e}^{\Omega W(x)} \left\{ \Omega^k \left[(W')^k b_k c\right] \right. \\
&\qquad \left. + \Omega^{k-1} \left[\frac{k(k-1)}{2} (W')^{k-2} W'' b_k c + k (W')^{k-1} (b_k c)'\right] + \mathcal{O}\left(\Omega^{k-2}\right) \right\}.
\end{aligned} \tag{79}$$

Substituting (79) into (74), at leading order $\mathcal{O}(\Omega)$ we recover equation (55) governing W, as discussed in the main text. At the next order $\mathcal{O}(1)$, we obtain the equation governing c, namely

$$0 = \sum_{k=2}^{\infty} \frac{(W')^{k-2}}{(k-2)!} \frac{W'' b_k c}{2} + \sum_{k=1}^{\infty} \frac{(W')^{k-1}}{(k-1)!} (b_k c)' \tag{80a}$$

$$\begin{aligned}
&= \frac{W'' b_2 c}{2} \sum_{k=0}^{\infty} \frac{(W')^{2k}}{(2k)!} + \frac{W'' b_1 c}{2} \sum_{k=1}^{\infty} \frac{(W')^{2k-1}}{(2k-1)!} \\
&\quad + (b_1 c)' \sum_{k=0}^{\infty} \frac{(W')^{2k}}{(2k)!} + (b_2 c)' \sum_{k=1}^{\infty} \frac{(W')^{2k-1}}{(2k-1)!}
\end{aligned} \tag{80b}$$

$$= \left[\frac{W'' b_2 c}{2} + (b_1 c)'\right] \cosh(W') + \left[\frac{W'' b_1 c}{2} + (b_2 c)'\right] \sinh(W'). \tag{80c}$$

Upon rearranging equation (80c) and integrating once, we obtain

$$c(x) = \frac{\hat{c}_\infty}{\sqrt{\bar{\lambda}(x)\bar{\mu}(x)}}, \tag{81}$$

for a normalising constant $\hat{c}_\infty$. Combining the results, the leading-order WKB approximation is given by

$$p_\infty(x) \sim \frac{\hat{c}_\infty \, \mathrm{e}^{\Omega\Phi(x)}}{\sqrt{\bar{\lambda}(x)\bar{\mu}(x)}}, \qquad \hat{c}_\infty \sim (1 - 1/\Lambda) \, \mathrm{e}^{-\Omega\Phi(x_e)} \sqrt{\frac{\Omega\Lambda}{2\pi}}, \tag{82}$$

where the constant $\hat{c}_\infty$ admits an asymptotic approximation via Laplace's method (see appendix A).

In principle, one could proceed to higher orders to determine asymptotic corrections to c in inverse powers of Ω.

Acknowledgements

The author would like to thank Prof. J. Frédéric Bonnans and Dr. Jakob Ruess for their insightful advice and helpful comments.

References

[1] M. Assaf, B. Meerson, and P. V. Sasorov. Large fluctuations in stochastic population dynamics: momentum-space calculations. *Journal of Statistical Mechanics: Theory and Experiment*, 2010(07):P07018, jul 2010.

[2] F. Baras, M. M. Mansour, and J. E. Pearson. Microscopic simulation of chemical bistability in homogeneous systems. *The Journal of Chemical Physics*, 105(18):8257–8261, 1996.

[3] G. E. P. Box. Science and statistics. *Journal of the American Statistical Association*, 71(356):791–799, 1976.

[4] P. C. Bressloff. *Stochastic Processes in Cell Biology*. Springer International Publishing, 2014.

[5] C. R. Doering, K. V. Sargsyan, and L. M. Sander. 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, 2005.

[6] S. Engblom. Computing the moments of high dimensional solutions of the master equation. *Applied Mathematics and Computation*, 180(2):498 – 515, 2006.

[7] A. Erdélyi. *Asymptotic expansions*. Dover Publications, 1956.

[8] C. Gardiner. *Stochastic Methods: A Handbook for the Natural and Social Sciences*. Springer Series in Synergetics. Springer Berlin Heidelberg, third edition, 2009.

[9] C. W. Gardiner and S. Chaturvedi. The poisson representation. I. a new technique for chemical master equations. *Journal of Statistical Physics*, 17(6):429–468, Dec 1977.

[10] T. S. Gardner, C. R. Cantor, and J. J. Collins. Construction of a genetic toggle switch in escherichia coli. *Nature*, 403(6767):339–342, 2000.

[11] M. A. Gibson and J. Bruck. Efficient exact stochastic simulation of chemical systems with many species and many channels. *The Journal of Physical Chemistry A*, 104(9):1876–1889, 2000.

[12] C. S. Gillespie. Moment-closure approximations for mass-action models. *IET Systems Biology*, 3(1):52–58, January 2009.

[13] D. T. Gillespie. The chemical Langevin equation. *The Journal of Chemical Physics*, 113(1):297–306, 2000.

[14] J. Grasman. The expected extinction time of a population within a system of interacting biological populations. *Bulletin of Mathematical Biology*, 58(3):555 – 568, 1996.

[15] R. Grima, P. Thomas, and A. V. Straube. How accurate are the nonlinear chemical Fokker-Planck and chemical Langevin equations? *J. Chem. Phys.*, 135(8):084103, 2011.

[16] E. J. Hinch. *Perturbation Methods*. Cambridge Texts in Applied Mathematics. Cambridge University Press, 1991.

[17] N. G. van Kampen. A power series expansion of the master equation. *Can. J. Phys.*, 39(4):551–567, 1961.

[18] S. G. Krantz and H. R. Parks. *A primer of real analytic functions*. Birkhäuser, 2 edition, 2002.

[19] B. J. Matkowsky, Z. Schuss, C. Knessl, C. Tier, and M. Mangel. First passage times for processes governed by master equations. In W. Horsthemke and D. K. Kondepudi, editors, *Fluctuations and Sensitivity in Nonequilibrium Systems*, pages 19–36, Berlin, Heidelberg, 1984. Springer Berlin Heidelberg.

[20] V. Méndez, M. Assaf, A. Masó-Puigdellosas, D. Campos, and W. Horsthemke. Demographic stochasticity and extinction in populations with allee effect. *Phys. Rev. E*, 99:022101, Feb 2019.

[21] R. Milo. What is the total number of protein molecules per cell volume? a call to rethink some published values. *BioEssays*, 35(12):1050–1055, 2013.

[22] L. Pauling. *General Chemistry*. Dover Books on Chemistry. Dover Publications, 1988.

[23] R. F. Pawula. Approximation of the linear Boltzmann equation by the Fokker–Planck equation. *Phys. Rev.*, 162:186–188, Oct 1967.

[24] H. Risken. *Fokker–Planck Equation*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1996.

[25] M. Strasser, F. J. Theis, and C. Marr. Stability and multiattractor dynamics of a toggle switch based on a two-stage model of stochastic gene expression. *Biophys. J.*, 102(1):19–29, 2012.

[26] P. Thomas, H. Matuschek, and R. Grima. How reliable is the linear noise approximation of gene regulatory networks? *BMC Genomics*, 14(4):S5, 2013.

[27] L. S. Tsimring. Noise in biology. *Reports on Progress in Physics*, 77(2):026601, jan 2014.

[28] N. G. van Kampen. *Stochastic processes in physics and chemistry*. North Holland, third edition, 2007.

[29] X. Yu and X.-Y. Li. Applications of wkb and fokker–planck methods in analyzing population extinction driven by weak demographic fluctuations. *Bulletin of Mathematical Biology*, 81(11):4840–4855, Nov 2019.