EPiC Series in Computing

Volume 114, 2026, Pages 10–37

Proceedings of the Workshop on Mathematical Modeling and Scientific Computing

Computing Extinction Barriers in a Quorum-Sensing Reaction Network

Mario Ayala and Johannes Zimmer

School of Computation, Information and Technology, Technische Universität München, Boltzmannstraße 3, 85748 Garching, Germany

mario.ayala@tum.de, jz@tum.de

Abstract

We introduce a simple reaction-network model of a quorum-sensing population that couples the cell density x to the signal density w. In the four-channel cell–signal network studied here, scaling signal production and removal by the same factor r leaves all deterministic equilibria, and their stability types, unchanged. Nevertheless, we show that r shifts the quasipotential barrier $\Delta V(r)$ for rare transitions towards extinction, and thus, under metastable exit assumptions, the mean time to reach a fixed neighbourhood of the extinction state on the exponential scale $e^{N\Delta V(r)}$.

We compute the barrier by minimization of the path action with the signal retained as a fluctuating coordinate, and we check it against mean extinction times obtained without the action solver, i.e., from the backward equation of the jump process on a truncated lattice and from exact stochastic simulation. Over a range of r, the minimum-action barriers satisfy

$$\Delta V(r) = \Delta V_\infty + O(1/r),$$

where ΔV_∞ is obtained by eliminating the signal first, and at $r = 1$ the barrier is 55% larger than ΔV_∞. A formal first-order expansion attributes this increase to the lag of the signal, of which the fluctuations of the signal remove a quarter. The saddle barrier $\Delta V(r)$ is also the least action needed to enter the basin of extinction, whereas the density threshold can be crossed more cheaply, by paths which keep the signal high and are usually followed by recovery. The mean extinction times of the backward equation, which include failed attempts, have a prefactor which tends to a constant, and they grow at a rate within 10^{-5} of $\Delta V(r)$ at every tested rate.

1 Introduction

Quorum sensing (QS) is a form of chemical communication in which cells release diffusible signalling molecules and respond to their concentration by altering gene expression. Among many others, QS regulates behaviours such as biofilm formation, exoenzyme production and bioluminescence, linking cellular activity to the surrounding population and its chemical environment [22, 29]. When these activities are costly to perform but benefit neighbouring cells, their regulation also shapes microbial cooperation [9].

The dynamics of this communication depend on both signal production and removal. Their speed introduces a timescale distinct from population growth and decline, raising the question of whether it can alter population persistence even when it leaves the deterministic equilibria in place. We address this question using a minimal, well-mixed stochastic model that explicitly retains both producing cells and the signal particles they release. More precisely, we investigate how this speed affects rare population extinction when the deterministic equilibrium locations and their stability types remain unchanged.

The model and the question

In this work, instead of starting from a population-level ODE description, we formulate our model from basic biological events, as is standard in stochastic population biology and reaction-network theory [2,14]. Every cell is a producer, and cell division, cell death, signal production and signal removal define a continuous-time Markov jump process for the cell and signal counts $n^C(t)$ and $n^w(t)$. At population scale N, the density process is

$$U^N(t) = \big(X^N(t), W^N(t)\big) := \left(\frac{n^C(t)}{N}, \frac{n^w(t)}{N}\right).$$

Throughout, we write elements of $\mathbb{R}^2$ as tuples, i.e., $U = (x, w)$, and identify them with column vectors, hence $\dot{U} = F(U)$ below is an equation between column vectors and, for $v \in \mathbb{R}^2$, $vv^\top$ is a 2×2 matrix. At a given density value (x, w), the increments and total rates of the four events are as follows.

event	**density increment**	**total rate**
cell division	$(1/N, 0)$	$Nx\, b(w)$
cell death	$(-1/N, 0)$	$Nx\, d(x)$
signal production	$(0, 1/N)$	$Nr\alpha_C x$
signal removal	$(0, -1/N)$	$Nr\kappa w$

Here $b(w)$ is the division rate, which increases with the signal through a Hill response, and $d(x)$ is the death rate, which includes a production cost and crowding. The constant α_C is the signal production rate per cell, κ is the signal removal coefficient, and $r > 0$ multiplies both signal channels. Hence r sets the speed of the signal dynamics while preserving the balance between production and removal at a fixed cell density. Raising r raises production while the cost c is held fixed, so varying r is a controlled mathematical experiment on the timescale of the signal, not an experimentally independent change of signal turnover; in a model where the cost scales with production, c and r would be coupled. Section 2 gives the closed forms of b and d, the generator, the biological rationale for the rates, and the dimensionless parameter values used below. These values are illustrative; our study is methodological rather than organism-specific.

One process, three regimes

At fixed r, the event rules above supply three complementary large-population descriptions: the law of large numbers (LLN), the central limit theorem (CLT) and the large-deviation principle

(LDP), given by

$$\begin{aligned} &\text{LLN:} && U^N \longrightarrow U, && \dot{U} = F(U), \\ &\text{CLT:} && \sqrt{N}\,(U^N - U) \Longrightarrow Z, && dZ = \nabla F(U)\, Z\, dt + D(U)^{1/2} dW, \\ &\text{LDP:} && \mathbb{P}(U^N \approx \phi) \asymp \exp\{-N\mathcal{S}_T(\phi)\}, && \mathcal{S}_T(\phi) = \int_0^T \mathbb{L}(\phi, \dot{\phi})\, dt. \end{aligned} \tag{1}$$

Here F is the mean-field drift (written out below), ∇F its linearization along the deterministic solution, D the diffusion matrix the reaction channels supply, W a standard planar Brownian motion, and $\mathbb{L}$ the Legendre transform of the limiting Hamiltonian. The LLN gives the deterministic trajectory, the CLT resolves fluctuations of order $N^{-1/2}$ around it, and the LDP assigns to a *rare* path ϕ a probability on the exponential scale $\exp\{-N\mathcal{S}_T(\phi)\}$ [11, 13].

The descriptions in (1) retain different information about the event rules. The deterministic drift combines opposing microscopic mechanisms through their net rates, whereas the fluctuation diffusion and path action also depend on the individual rates. Hence neither the deterministic drift nor its equilibria, described below, determine the cost of extinction.

In order to explain this in more detail, we start from the LLN. By Kurtz's theorem [20] (Theorem 1 in this work), U^N converges to a deterministic limit $U = (x, w)$ satisfying the following ODE system

$$\dot{U} = F(U) = \begin{pmatrix} x\big[b(w) - d(x)\big] \\ r\big[\alpha_C x - \kappa w\big] \end{pmatrix}.$$

Therefore, at equilibrium, the signal satisfies

$$w = \bar{w}(x) := \frac{\alpha_C x}{\kappa},$$

and every positive equilibrium solves $b(\bar{w}(x)) = d(x)$, independently of r. For the bistable parameters studied here, the system has three equilibria: the extinction state $U_{\text{off}} = (0, 0)$, a positive saddle state $U_* = (x^*, \bar{w}(x^*))$, and a stable cooperative state $U_{\text{on}} = (x_{\text{on}}, \bar{w}(x_{\text{on}}))$. Proposition 2 shows that all three equilibria retain their stability types as r varies. Nevertheless, their relaxation rates, trajectories and basins need not remain unchanged.

Extinction

In this work, by population extinction we mean a transition from the stable cooperative state U_{on} towards U_{off}. The deterministic attraction basins of these states are separated by a curve, the separatrix, whose location in general depends on r, even though the equilibria themselves do not. The transition therefore involves the joint cell-signal state and cannot be identified by a cell-density threshold alone.

In order to measure this transition, we choose a fixed small neighbourhood R_{off} (see (8)) of the extinction state U_{off}, contained in its deterministic attraction basin, and define, for the process started near U_{on},

$$\tau_N^{\text{off}}(r) := \inf\{t \geq 0 : U^N(t) \in R_{\text{off}}\}, \tag{2}$$

i.e., the first arrival in a region from which the deterministic evolution tends to U_{off}. We call $\tau_N^{\text{off}}(r)$ the extinction time. Entry into R_{off} is not yet cell extinction, $X^N = 0$, and does not exclude subsequent stochastic recovery. We quantify persistence by means of the mean of this stopping time and its dependence on the population scale N. The target set is specified in more

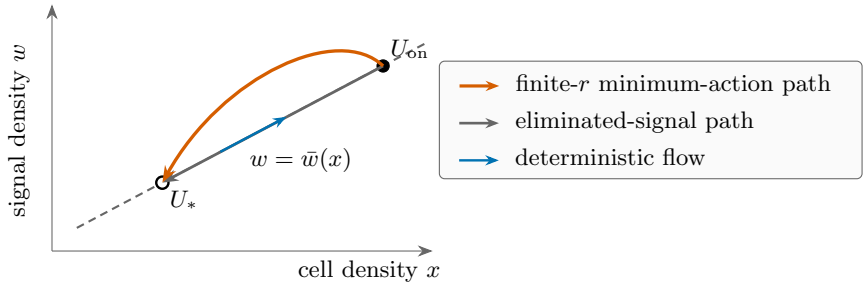


Figure 1: Why signal elimination changes the barrier without changing the deterministic equilibria. Both paths run from U_{on} to the saddle U_*, the endpoint of the barrier (Section 4). Elimination confines the path to the dashed relation $w = \bar{w}(x)$; at finite r the signal lags above it, so the minimum-action path joins the same endpoints through different states.

detail in Section 3.1.2. The justification for using the saddle action as the transition barrier, and the assumptions needed to connect that barrier to mean extinction times, are given in Section 4. Section 5.2 states which stopping times are computed.

For large r, it is natural at the deterministic level to replace the signal by $\bar{w}(x)$. Whether this reduction also preserves the probability of rare extinction is a separate question. At finite r, a fluctuating path towards extinction can explore the joint cell-signal plane: the signal need not remain at its quasi-steady value while the cell density falls. The quantity we compute measures the cost of this rare transition from the cooperative state U_{on} to the saddle U_*, which is the least-cost entrance to the basin of extinction (Section 4). Figure 1 sketches the basic mechanism.

We denote the least action from U_{on} to U_* by $\Delta V(r)$, the quasipotential barrier; its precise definition is given in Section 4. Subject to metastable exit estimates for the jump process, which we state but do not prove, this barrier governs the mean of the extinction time (2) on the exponential scale

$$\mathbb{E}[\tau_N^{\text{off}}(r)] \asymp \exp\{N\Delta V(r)\}.$$

We compare it with ΔV_∞, the barrier obtained by eliminating the signal before constructing the reduced path action.

Contribution and main findings

The limit theorems and the large-deviation principle are established tools. Our contribution is a calibrated barrier computation for this four-channel network, checked independently: we minimize the path action with the signal retained as a fluctuating coordinate, show that the saddle barrier is the least interior action needed to enter the basin of extinction (Proposition 7), and compare with mean extinction times obtained from the backward equation of the jump process and from exact stochastic simulation, neither of which uses the action solver. The minimum-action results are compared with the expansion

$$\Delta V(r) = \Delta V_\infty + \frac{A}{r} + O(r^{-2}), \qquad r \to \infty, \tag{3}$$

whose coefficient A is the prediction of a formal boundary-layer calculation. We sketch this calculation in Section 4.2 and refer to [6] for the details; we do not prove the expansion here. Over the values of r studied, the minimum-action results support the $1/r$ correction, and at $c = 0.36$ and $r = 1$ the computed barrier is 55% larger than ΔV_∞. The same calculation

separates the lag of the signal, which raises the barrier, from the fluctuations of the signal, which remove a quarter of this increase.

For the extinction time we solve the backward equation of the jump process on a truncated lattice, for $40 \le N \le 400$. The prefactor $\mathbb{E}[\tau_N^{\mathrm{off}}(r)]\, e^{-N\Delta V(r)}$ tends to a constant, with a correction of order $1/N$, and the fitted growth rate differs from the minimum-action barrier by less than 10^{-5} at each of the five rates. This separates it from ΔV_∞ at these rates, and the simulated means agree with the exact ones within sampling error. The density threshold x^* behaves differently: at $r = 0.5$ a threshold point with higher signal costs 7.7% less action than the saddle, and in simulation most threshold crossings at this rate are followed by recovery. Since the barrier enters the mean extinction time multiplied by N, an r-dependent barrier changes the mean extinction time by a factor exponential in N, although the equilibria do not move.

That the elimination of a fast variable can change the rate of a rare event is well known. Large deviations of fast-slow systems are analysed in [7], the elimination of fast stochastic variables at the level of the linear noise approximation in [27], and the breakdown of fast-slow reductions for escape rates in [23]. In population dynamics, WKB methods for extinction are reviewed in [3], and [19] shows that an extinction rate can change exponentially under a small change of the parameters. The closest setting to ours is the genetic switch of [4], where the fast species is the mRNA: there the mRNA noise changes the action at leading order, through bursts of protein production, and the correction in the ratio of the two lifetimes is stated as an order of magnitude. In the network we study here, the fast variable is as abundant as the slow one, hence the leading order is the eliminated-signal barrier ΔV_∞ and the effect of the signal starts at order $1/r$. The minimum action method and the backward equation are standard tools. Relative to these works, the present paper gives a quantified example in a quorum-sensing network: the barriers at finite r with an independent check by the backward equation, the numerical test of the coefficient A, of its separation into lag and noise and of the range of r in which the $1/r$ law holds, and the comparison between the threshold, the recovery and the basin pictures of extinction. We are not aware of a formula for the first-order coefficient in the cited works. The formula (22) is obtained in the forthcoming paper [6], and its derivation is not claimed as a contribution of this paper.

Outline

Section 2 gives the reaction-network formulation and parameters. Section 3 develops the deterministic dynamics, Gaussian fluctuations and large deviations. Section 4 defines the extinction barrier, relates it to the basin of U_{off}, and obtains the eliminated-signal benchmark ΔV_∞. Section 5 presents the action minimization, the backward equation and the simulations, Section 6 compares their results, and Section 7 discusses the conclusions and remaining limitations.

2 The reaction network

2.1 Events, state, and generator

We now describe in detail the reaction network introduced in Section 1. Signal increases the cell division rate through a sigmoidal response, production contributes a private cost to cell death, and crowding limits growth. Particles are created and destroyed one at a time, and no event changes both populations at the same time. The factor r multiplies production and removal

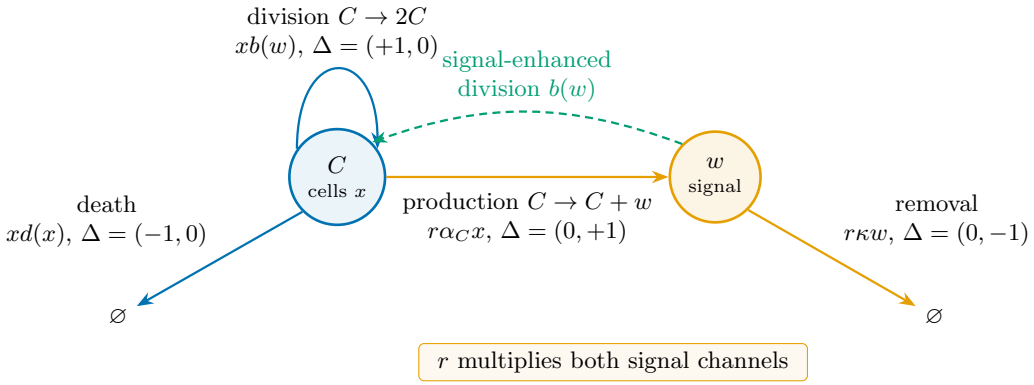


Figure 2: The four stochastic channels and their feedback. Blue arrows change the cell count and orange arrows change the signal count. The dashed arrow is regulation, not an additional jump: signal enters the division rate $b(w)$.

together, so it changes the speed of the signal dynamics without changing the stationary signal level (Proposition 2).

The state is the density $U = (x, w)$ of the introduction, with $x = n^C/N$ and $w = n^w/N$, and the system is assumed well mixed throughout. Each of the four distinct mechanisms changes exactly one coordinate of U by $1/N$, while the total event rate is of order N. The resulting density process has a deterministic large-population limit and rare-event probabilities on the exponential scale N. Notice that cells and signal share the population scale N, i.e., the signal is as abundant as the cells. This is a modelling choice, and Section 6 states what changes when the signal count is scaled as νN instead.

We model these mechanisms as the reaction channels of a stochastic reaction network [2]: a channel ϱ is one reaction, specified by a fixed increment Δ_ϱ and a rate density β_ϱ. The time-evolution of the state U is therefore a continuous-time Markov jump process with state space $(N^{-1}\mathbb{N}_0) \times (N^{-1}\mathbb{N}_0)$: each channel carries an exponential clock of rate $N\beta_\varrho(U)$ depending on the current state U, and its firing moves the state by a fixed increment

$$U \longmapsto U + \frac{\Delta_\varrho}{N}.$$

The channels, their increments, and their rate densities are summarized in Fig. 2.

Notice that our model is microscopic: it specifies which events occur and at what rate, and it imposes nothing at the population level. The dynamics of U are best described in terms of its infinitesimal generator, which acts on functions of the density state by

$$(\mathfrak{L}_N f)(U) = \sum_\varrho N\beta_\varrho(U)\left[f\left(U + \frac{\Delta_\varrho}{N}\right) - f(U)\right], \tag{4}$$

i.e., the time-evolution of U is a density-dependent Markov chain in the sense of Kurtz [12, 20]. Everything computed below is obtained from (4): the large-population descriptions of Section 3 are consequences of it rather than starting points, and so is the barrier of Section 4.

2.2 Rates and parameters

Each of the three functional forms below is chosen for a reason rather than for tractability. The division response is of Hill type because quorum-sensing regulation is cooperative: a receptor

Table 1: Parameters of the four-channel reaction network; the increments and rates of the channels are those of the table in the introduction. The cost c is varied where stated, with $c = 0.36$ as reference, and r is swept over two decades. All quantities are dimensionless.

channel ϱ	symbol	meaning	value
division	b_0	basal division rate	0.5
	b_1	maximal signal-induced increment	2.0
	K_h	half-response signal level	1.0
	h	Hill exponent (cooperativity)	4
death	d_0	basal death rate	0.8
	c	cost of production	0.36
	θ	crowding coefficient	1.0
production	α_C	production rate per cell	2.0
removal	κ	signal removal rate	1.0
signal channels	r	common factor of production and removal	varied

of LuxR type binds the signal and the complex multimerizes before it binds the promoter, so the response to signal is switch-like rather than graded, and the Hill exponent h is what carries that cooperativity [17,22,29]. Production enters the death rate because making signal is metabolically costly, which is also what makes the behaviour exploitable by non-producers [9]. The remaining term in the death rate is ordinary logistic self-limitation. Hence

$$b(q) = b_0 + b_1 \frac{q^h}{K_h^h + q^h}, \qquad d(x) = d_0 + c + \theta x. \tag{5}$$

The sigmoidal response in (5) permits bistability for the dimensionless values of Table 1, which are used throughout. They are illustrative, hence the conclusions below are methodological rather than organism-specific.

3 Three regimes

The large N behaviour of the process admits three descriptions, distinguished by the size of the deviation from the deterministic path that each one resolves: a law of large numbers for the path itself, a central limit theorem for deviations of order $N^{-1/2}$, and a large-deviation principle for deviations of order one. All three are known results in the context of reaction networks. We state them for the network of Fig. 2 and verify their hypotheses, and we cite rather than reproduce their proofs. The barrier computed in this paper is an object of the third regime, and (1) of the introduction states all three side by side.

3.1 Law of large numbers

Write $F(U) = \sum_\varrho \Delta_\varrho \beta_\varrho(U)$ for the drift associated with (4); summing the four increments of Fig. 2 against their rate densities gives the vector field of the introduction.

Theorem 1 (Kurtz [20]; see also [12, Ch. 11])**.** *Fix $r > 0$, let $G \subset (0,\infty)^2$ be open and bounded, and suppose $U^N(0) \to U_0 \in G$ in probability. Let $U = (x, w)$ solve $\dot U = F(U)$, i.e.,*

$$\dot x = x\big[b(w) - d(x)\big], \qquad \dot w = r\big[\alpha_C x - \kappa w\big], \tag{6}$$

with $U(0) = U_0$, and let $T > 0$ be such that $U(t) \in G$ for every $t \le T$. Then for every $\varepsilon > 0$,

$$\lim_{N\to\infty} \mathbb{P}\Big(\sup_{t\le T} |U^N(t) - U(t)| > \varepsilon\Big) = 0.$$

The application of [20] is immediate here: the family is finite, the increments Δ_ϱ have unit length, and, since h is a positive integer and $K_h^h + w^h$ does not vanish, every rate density β_ϱ of Fig. 2 is smooth on $[0,\infty)^2$ and therefore Lipschitz on compact subsets of it. Notice that G is required to be bounded because $xd(x)$ grows quadratically; the restriction is harmless below, where every trajectory of interest stays in a compact subset of $(0,\infty)^2$.

3.1.1 Equilibria and stability

Within this mean-field system, the formal limit $r \to \infty$ slaves the signal to

$$\bar{w}(x) = \frac{\alpha_C x}{\kappa}. \tag{7}$$

The order of this statement matters. At fixed N, the fast signal has a conditional Poisson equilibrium rather than the deterministic value (7); deterministic slaving is the fast limit after the mean-field reduction. In this subsection, however, $\bar{w}$ serves only to locate and classify the equilibria of the full system (6), not to reduce it; the dynamics with the signal eliminated are taken up in Section 4.1. Writing $\bar{b}(x) := b(\bar{w}(x))$ for the division rate along (7) and $\bar{F}(x) := x[\bar{b}(x) - d(x)]$ for the slaved drift, the factor r leaves the deterministic picture untouched.

Proposition 2. *Let $\theta > 0$ and let $U_{\mathrm{eq}} = (x_{\mathrm{eq}}, w_{\mathrm{eq}})$ be an equilibrium of (6). Then $w_{\mathrm{eq}} = \bar{w}(x_{\mathrm{eq}})$ and $\bar{F}(x_{\mathrm{eq}}) = 0$, both independently of r. Moreover, the Jacobian J_r of (6) at U_{eq} satisfies*

(i) If $x_{\mathrm{eq}} > 0$, then

$$\det J_r = -r\kappa\, \bar{F}'(x_{\mathrm{eq}}), \qquad \operatorname{tr} J_r = -x_{\mathrm{eq}} d'(x_{\mathrm{eq}}) - r\kappa < 0,$$

hence for every $r > 0$ the equilibrium is a saddle when $\bar{F}'(x_{\mathrm{eq}}) > 0$ and asymptotically stable when $\bar{F}'(x_{\mathrm{eq}}) < 0$.

(ii) If $x_{\mathrm{eq}} = 0$, then $U_{\mathrm{eq}} = U_{\mathrm{off}} := (0,0)$ and J_r has eigenvalues $b_0 - d_0 - c$ and $-r\kappa$, hence for every $r > 0$ the extinction state is asymptotically stable when $b_0 < d_0 + c$ and a saddle when $b_0 > d_0 + c$.

Proof. The second component of (6) vanishes if and only if $w = \bar{w}(x)$, and the first then reads $\bar{F}(x) = 0$; neither involves r. (i) Using $\bar{b}(x_{\mathrm{eq}}) = d(x_{\mathrm{eq}})$,

$$J_r = \begin{pmatrix} -x_{\mathrm{eq}} d'(x_{\mathrm{eq}}) & x_{\mathrm{eq}} b'(w_{\mathrm{eq}}) \\ r\alpha_C & -r\kappa \end{pmatrix}, \qquad \bar{F}'(x_{\mathrm{eq}}) = x_{\mathrm{eq}}\big[(\alpha_C/\kappa) b'(w_{\mathrm{eq}}) - d'(x_{\mathrm{eq}})\big],$$

whence $\det J_r = -r\kappa \bar{F}'(x_{\mathrm{eq}})$, and the trace is negative since $d' \equiv \theta > 0$. (ii) At the origin J_r is lower triangular with diagonal entries $b_0 - d_0 - c$ and $-r\kappa$. □

For Table 1 at $c = 0.36$ the system has the three equilibria of Fig. 3: the stable extinction state U_{off}, an unstable threshold x^*, and a stable cooperative state x_{on}, with

$$x^* = 0.55881, \qquad x_{\mathrm{on}} = 1.29675, \qquad \bar{w}(x_{\mathrm{on}}) = 2.59350.$$

Since $b_0 < d_0$ in Table 1, U_{off} is asymptotically stable for every cost $c \ge 0$ and every $r > 0$.

3.1.2 Basins of attraction and the extinction target

Let $\Phi_t^{(r)}(u)$ denote the deterministic solution starting from u, and write $Q = (0,\infty)^2$. The two attraction basins in the open quadrant are

$$\mathcal{B}_{\mathrm{on}}(r) := \{u \in Q : \Phi_t^{(r)}(u) \to U_{\mathrm{on}}\}, \qquad \mathcal{B}_{\mathrm{off}}(r) := \{u \in Q : \Phi_t^{(r)}(u) \to U_{\mathrm{off}}\}.$$

For Table 1 and each cost $c \in [0.2, 0.5]$, which contains the costs used below, $b_0 < d_0 + c < b_0 + b_1$. The positive zeros of the slaved drift $\bar{F}$ are the positive roots of

$$P(x) := \big(b_0 - d_0 - c - \theta x\big)\big(K_h^h + (\alpha_C x/\kappa)^h\big) + b_1(\alpha_C x/\kappa)^h,$$

obtained by multiplying $b(\bar{w}(x)) = d(x)$ by the positive factor $K_h^h + (\alpha_C x/\kappa)^h$. For $h = 4$ the polynomial P has degree 5, and its nonzero coefficients, in the degrees $0, 1, 4, 5$, have the signs $-, -, +, -$. By Descartes' rule of signs P has at most two positive roots, counted with multiplicity. Since $P(0) < 0$, $P(1) = 9.9 - 17c > 0$ and $P(x) \to -\infty$ as $x \to \infty$, it has exactly two, $x^* < x_{\mathrm{on}}$, and both are simple. The equilibria in $[0,\infty)^2$ are then exactly U_{off}, the saddle $U_* = (x^*, w^*)$, where $w^* := \bar{w}(x^*)$, and U_{on}. Since $\bar{F} < 0$ on $(0, x^*)$ and both zeros are simple, $\bar{F}'(x^*) > 0 > \bar{F}'(x_{\mathrm{on}})$, hence by Proposition 2, for every $r > 0$, U_* is a hyperbolic saddle and U_{on}, U_{off} are asymptotically stable.

The two basins are separated by the stable manifold $W^s(U_*)$ of the saddle, the separatrix, i.e., the set of initial states whose solutions converge to U_*. Although the equilibria are independent of r, the separatrix generally is not (Fig. 3). Thus the cell density alone does not determine which equilibrium a trajectory approaches. Two pieces of the picture are nevertheless fixed for all r. One can show that, for every $r > 0$, the quadrant $\{(x, w) : x > x^*,\ w > w^*\}$ lies in $\mathcal{B}_{\mathrm{on}}(r)$ and the segment $\{x^*\} \times (0, w^*)$ lies in $\mathcal{B}_{\mathrm{off}}(r)$; we omit the proof. In particular, the threshold line $\{x = x^*\}$ meets the separatrix only at U_*, and a crossing of that line with $w > w^*$ lies in the cooperative basin.

To measure a transition towards extinction, we record arrival in a specified neighbourhood of U_{off} within its attraction basin. This places the stopping target beyond the separatrix, rather than at the first crossing of it. Specifically, choose

$$R_{\mathrm{off}} := [0, a_{\mathrm{off}}] \times [0, b_{\mathrm{off}}], \qquad 0 < a_{\mathrm{off}} < x^*, \qquad b(b_{\mathrm{off}}) < d_0 + c, \qquad \alpha_C a_{\mathrm{off}} < \kappa b_{\mathrm{off}}. \tag{8}$$

Such a_{off}, b_{off} exist whenever $b_0 < d_0 + c$, by continuity of b.

Proposition 3. *Let $a_{\mathrm{off}}, b_{\mathrm{off}} > 0$ satisfy the last two conditions in (8). Then, for every $r > 0$, R_{off} is forward invariant under (6) and every solution started in R_{off} converges to U_{off}. In particular $R_{\mathrm{off}} \cap Q \subset \mathcal{B}_{\mathrm{off}}(r)$.*

Proof. Let $\gamma := d_0 + c - b(b_{\mathrm{off}}) > 0$. On R_{off}, $\dot{x} = x[b(w) - d(x)] \le -\gamma x$, and on the edge $w = b_{\mathrm{off}}$, $\dot{w} = r(\alpha_C x - \kappa b_{\mathrm{off}}) \le r(\alpha_C a_{\mathrm{off}} - \kappa b_{\mathrm{off}}) < 0$. Hence the flow points into R_{off} on the edges $x = a_{\mathrm{off}}$ and $w = b_{\mathrm{off}}$. Moreover, the axis $\{x = 0\}$ is invariant, and on the edge $\{w = 0,\ x > 0\}$ we have $\dot{w} = r\alpha_C x > 0$, i.e., the flow points inward there as well, so R_{off} is forward invariant. Inside it $x(t) \le x(0)e^{-\gamma t}$, and the variation-of-constants formula for w gives, with $m := \min\{r\kappa, \gamma\}$,

$$w(t) = e^{-r\kappa t}w(0) + r\alpha_C \int_0^t e^{-r\kappa(t-s)}x(s)\,ds \le e^{-r\kappa t}w(0) + r\alpha_C x(0)\,t\,e^{-mt} \to 0.$$

□

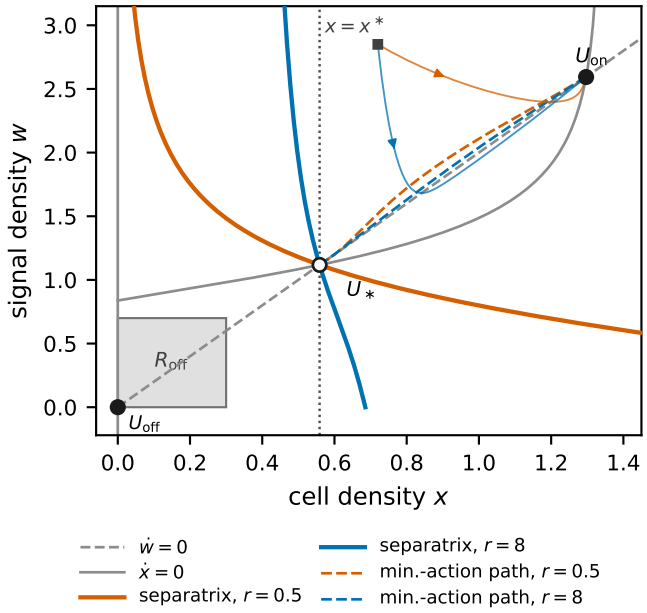


Figure 3: Phase portrait of (6) at $c = 0.36$. Grey: the nullclines $\dot{w} = 0$ (dashed) and $\dot{x} = 0$ (solid), which do not depend on r; filled points are the stable states and the open point is the saddle U_*. Thick curves: the separatrix at $r = 0.5$ and $r = 8$, with $\mathcal{B}_{\text{off}}(r)$ below and to the left of it. Dotted: the threshold line $x = x^*$. Shaded: the target $R_{\text{off}} = [0, 0.3] \times [0, 0.7]$. Dashed coloured curves: minimum-action paths from U_{on} to U_*. Thin curves: relaxation paths from the square marker, in the same colours, which reach the same state along routes that depend on r. All quantities are dimensionless.

The first condition in (8) keeps the target at a positive distance from the saddle.

Starting near U_{on}, the extinction time (2) is the first entry of U^N into this set, with the target fixed independently of N and r. Its deterministic invariance does not make it stochastically absorbing: recovery remains possible after entry. This event differs from both the population-threshold crossing $X^N \leq x^*$ and exact cell extinction, $X^N = 0$. The relation between the action required to reach this target and the saddle action is established in Section 4.

3.2 Gaussian fluctuations

We are now interested in the fluctuations of U^N about the deterministic solution U of (6). Theorem 1 makes their difference vanish, hence we magnify it and define the fluctuation field

$$Z^N := \sqrt{N}\,(U^N - U), \tag{9}$$

$\sqrt{N}$ being the scale at which it has a nondegenerate limit. Deviations of that order about the deterministic path are Gaussian, provided the initial fluctuation is.

Theorem 4 (Kurtz; see [12, Ch. 11])**.** *Assume the hypotheses of Theorem 1 and let $Z^N(0) \Rightarrow Z_0$ for a random vector Z_0 that is Gaussian or deterministic. Then, on the time interval of Theorem 1, the field (9) converges weakly in the Skorokhod space $D([0,T];\mathbb{R}^2)$, $Z^N \Rightarrow Z =$*

(Z_x, Z_w), where Z is the Gaussian process with $Z(0) = Z_0$ solving

$$\begin{aligned} dZ_x &= \big[\big(b(w) - d(x) - x\,d'(x)\big)Z_x + x\,b'(w)\,Z_w\big]dt + \sqrt{x\,b(w)}\;dW_1 - \sqrt{x\,d(x)}\;dW_2, \\ dZ_w &= r\big[\alpha_C Z_x - \kappa Z_w\big]dt + \sqrt{r\alpha_C x}\;dW_3 - \sqrt{r\kappa w}\;dW_4, \end{aligned} \tag{10}$$

where $(x, w) = U(t)$ is the deterministic solution of (6) and $W_1, \dots, W_4$ are independent standard Brownian motions, one per channel of Fig. 2, independent of Z_0.

The cited theorem asks one more derivative of the rate densities and a moment condition of the jumps, both available since every β_ϱ is smooth on $[0,\infty)^2$ and the increments have unit length. Collecting the four noise terms gives the diffusion matrix $D = \sum_\varrho \Delta_\varrho \Delta_\varrho^\top \beta_\varrho$, in which, unlike the drift, the increments enter quadratically. The channels of Fig. 2 are axis-aligned, hence

$$D(x, w) = \operatorname{diag}\big(xb(w) + xd(x),\; r[\alpha_C x + \kappa w]\big). \tag{11}$$

At an asymptotically stable equilibrium the stationary covariance Σ of Z solves the Lyapunov equation $J_r\Sigma + \Sigma J_r^\top + D = 0$. At U_{on} the variance of the signal conditional on the cells is $1.88\,\bar{w}(x_{\mathrm{on}})$ at $r = 0.5$, $2.01\,\bar{w}(x_{\mathrm{on}})$ at $r = 1$ and $1.02\,\bar{w}(x_{\mathrm{on}})$ at $r = 256$, i.e., this regime recovers the conditional Poisson equilibrium at large r and departs from it by a factor near two for $r \le 2$.

3.3 Large deviations

By means of the exponential change of variable applied to (4), i.e., of the nonlinear generator $N^{-1}e^{-Nf}\mathfrak{L}_N(e^{Nf})$ of the Feng-Kurtz construction [13], we obtain in the limit the Hamiltonian

$$\mathbb{H}(U, p) = \sum_\varrho \beta_\varrho(U)\big(e^{\langle p, \Delta_\varrho\rangle} - 1\big),$$

which for the four channels of Fig. 2 and momenta (p, q) reads

$$\mathbb{H}(x, w, p, q) = \underbrace{xb(w)(e^{p} - 1) + xd(x)(e^{-p} - 1)}_{\mathbb{H}_s} + r\underbrace{\big[\alpha_C x(e^{q} - 1) + \kappa w(e^{-q} - 1)\big]}_{\mathbb{H}_f}. \tag{12}$$

Its Legendre-Fenchel transform $\mathbb{L}(U, v) = \sup_p[\langle p, v\rangle - \mathbb{H}(U, p)]$ is the local rate, and the action of an absolutely continuous path is $\mathcal{S}_T(U) = \int_0^T \mathbb{L}(U(t), \dot{U}(t))\,dt$. Deviations of order one are exponentially unlikely at speed N.

Theorem 5 (Agazzi et al. [1, Thm. 1 and Rem. 2.6]; see also [26])**.** *Let G be an open set whose closure is a compact subset of $Q = (0,\infty)^2$, and let the initial states be deterministic, $U^N(0) = u_N \to u_0 \in G$. Let $\widetilde{U}^N$ be the process with the jumps of Fig. 2 and rate densities $\widetilde{\beta}_\varrho$ that coincide with β_ϱ on $\overline{G}$ and are Lipschitz, bounded and bounded away from zero on $\mathbb{R}^2$, started at $\widetilde{U}^N(0) = u_N$. Then $\widetilde{U}^N$ satisfies a large-deviation principle on $[0, T]$ at speed N with the good rate function*

$$I_{u_0}(U) = \begin{cases} \widetilde{\mathcal{S}}_T(U), & U \in \mathrm{AC}([0,T];\mathbb{R}^2),\ U(0) = u_0, \\ +\infty, & \text{otherwise,} \end{cases}$$

where $\widetilde{\mathcal{S}}_T$ is the action built from $\widetilde{\beta}_\varrho$. That is, the sublevel sets of I_{u_0} are compact, and for every closed set C and every open set O of the Skorokhod space $D([0,T];\mathbb{R}^2)$ with the J_1 topology,

$$\limsup_{N\to\infty} \frac{1}{N} \log \mathbb{P}\big(\widetilde{U}^N \in C\big) \le - \inf_{U\in C} I_{u_0}(U),$$
$$\liminf_{N\to\infty} \frac{1}{N} \log \mathbb{P}\big(\widetilde{U}^N \in O\big) \ge - \inf_{U\in O} I_{u_0}(U).$$

Such extensions exist, e.g. the rate densities composed with the projection onto a compact rectangle in Q containing $\overline{G}$. The processes can be coupled up to their first exit from G. As a consequence, the displayed bounds also hold for U^N for open or closed path events all of whose paths remain in G, with $\widetilde{\mathcal{S}}_T$ replaced by $\mathcal{S}_T$. This is the only localization used below. Results of this type go back to [21,26], and the flux formulation is given in [24]. The localization is needed because the rate densities vanish on the axes. The cited theorem also covers such degenerate sets, under an escape condition on the rates near them that we do not verify for this network. The variational results of Section 4 restrict every admissible path to the open quadrant. The local formulation above does not by itself give an exit-time theorem for the original chain, nor a comparison with paths that touch the axes.

Remark 6. The third regime contains the other two, in the sense that they are the first two terms of $\mathbb{H}$ at zero momentum. Differentiating $\mathbb{H}(U,p) = \sum_\varrho \beta_\varrho(U)(e^{\langle p,\Delta_\varrho\rangle} - 1)$ once and twice at $p = 0$ gives $\sum_\varrho \beta_\varrho \Delta_\varrho = F(U)$ and $\sum_\varrho \beta_\varrho \Delta_\varrho \Delta_\varrho^\top = D(U)$, i.e., the drift of (6) and the diffusion matrix (11), hence

$$\mathbb{H}(U,p) = \langle F(U), p\rangle + \tfrac{1}{2}\langle p, D(U)p\rangle + O(|p|^3),$$
$$\mathbb{L}(U,v) = \tfrac{1}{2}\langle v - F(U), D(U)^{-1}\big(v - F(U)\big)\rangle + O(|v - F(U)|^3).$$

The cost of a velocity is therefore measured, to leading order, in the norm that the mobility of Theorem 4 supplies, and to that order the action is the one of the diffusion driven by the noise of (10). Notice that this is also why the calibration of Section 5.1 works: near U_{on} the quasipotential is governed by the quadratic form above, whose metric is D^{-1}. The terms of third and higher order are where the increments enter exponentially rather than averaged as in F or quadratically as in D. At the parameters of this paper they are a small correction to the barrier: on the eliminated-signal instanton $|p| \le 0.131$, the barrier of the quadratic Hamiltonian differs from ΔV_∞ by 0.1%, and Section 6 gives the same comparison at finite r. Two networks with the same drift and different channels therefore share Theorem 1 and can differ in both later regimes. The drift carries net rates and the diffusion gross ones: at U_{on} the cell drift vanishes because $b = d$, whereas the cell diffusion is $2x\,b(w)$. The factor r is a case in point. By Proposition 2 it leaves the equilibria and their stability types unchanged, but it survives in J_r and in D_{ww}, hence in the stationary covariance Σ, and in the barrier. Since the rate densities vanish simultaneously only at the origin, U_{off} is the only absorbing state of U^N, and the cooperative state is metastable for the jump process: its location is common to every r, its spread Σ/N depends on r, and, subject to the exit estimates discussed after (15), its lifetime is of order $\exp\{N\Delta V(r)\}$.

4 The extinction barrier

The extinction time (2) is the first entry into the target R_{off} of Section 3.1.2. This section relates the least action needed to reach that target to the least action needed to reach the

saddle U_*, which is the quantity we compute. Theorem 5 assigns the action $\mathcal{S}_T(U)$ to an absolutely continuous path U with values in a compact subset of Q. On the exponential scale, a small neighbourhood of such a path has probability of order $\exp\{-N\mathcal{S}_T(U)\}$, so the likeliest ways of reaching a target are paths of nearly minimal action. Accordingly, we restrict every admissible path to the open quadrant and define, for $u, z \in Q$,

$$V_r^Q(u,z) := \inf_{T>0} \inf_{\substack{U\in \mathrm{AC}([0,T];Q):\\ U(0)=u,\ U(T)=z}} \mathcal{S}_T(U), \qquad V_r^Q(z) := V_r^Q(U_{\mathrm{on}}, z).$$

The superscript records the restriction to Q. The quasipotential barrier is the value at the saddle

$$\Delta V(r) := V_r^Q(U_*) = \inf_{T>0} \inf_{\substack{U\in \mathrm{AC}([0,T];Q):\\ U(0)=U_{\mathrm{on}},\ U(T)=U_*}} \mathcal{S}_T(U). \tag{13}$$

Here the first infimum allows the path to choose its duration. This is needed because the minimizing orbit approaches the two fixed points only asymptotically. Hence no finite value of T realizes the free-time minimum; instead, the minimum is approached as $T \to \infty$. This is the reason for the horizon study in Section 5.1. The stored minimizers at $c = 0.36$, $r \in \{0.5, 1, 4, 16\}$, stay in $[x^*, \infty)^2$, well inside Q.

Equation (13) is a genuine minimization problem. The local rate is nonnegative and vanishes along the drift, i.e., $\mathbb{L}(U, F(U)) = 0$, since $\mathbb{H}(U, 0) = 0$, $p \mapsto \mathbb{H}(U, p)$ is convex and $\nabla_p \mathbb{H}(U, 0) = F(U)$ (Remark 6); hence following the deterministic drift costs nothing. Leaving the basin of the cooperative equilibrium requires a fluctuation; after entering the off basin, the deterministic evolution can carry the system towards the off state at zero action. Moreover, at finite r the state is two-dimensional: the signal need not remain at $w = \bar{w}(x)$. The cheapest path can exploit signal lag, as illustrated in Fig. 1. Since the network is not reversible, there is also no general reason for this path to be the deterministic relaxation path run backwards.

The saddle enters as in Freidlin-Wentzell theory [15]. A path from U_{on} into the off basin must cross the separatrix, and from any point of the separatrix the deterministic flow leads to U_* at zero action, so no point of the separatrix is cheaper to reach than U_*. Conversely, from U_* the branch of the unstable manifold that leaves towards smaller x converges to U_{off}, hence enters R_{off}, at zero action.

Proposition 7. *Let $c \in [0.2, 0.5]$ and $r > 0$, and let $R \subset [0,\infty)^2$ be a relative neighbourhood of U_{off} with $R \cap Q \subset \mathcal{B}_{\mathrm{off}}(r)$. Then $\inf_{z\in R\cap Q} V_r^Q(z) = \Delta V(r)$. In particular, by Proposition 3,*

$$\inf_{z\in R_{\mathrm{off}}\cap Q} V_r^Q(z) = \Delta V(r). \tag{14}$$

The proof is given in Appendix A. Its lower-bound step uses only that the endpoint lies outside $\mathcal{B}_{\mathrm{on}}(r)$, hence it also gives $V_r^Q(z) \geq \Delta V(r)$ for every $z \in Q \setminus \mathcal{B}_{\mathrm{on}}(r)$. The proposition concerns an infimum of action and does not compare individual paths. Reaching U_* does not force extinction, since noise near the saddle can send the process to either basin.

The identification (14) concerns actions. Relating it to the mean of the extinction time (2) additionally requires metastable exit estimates for the jump process, appropriate to the target R_{off} and to the vanishing of the rates on the coordinate axes; we do not derive them here. Subject to such estimates, the prediction is

$$\lim_{N\to\infty} \frac{1}{N} \log \mathbb{E}\big[\tau_N^{\mathrm{off}}(r)\big] = \Delta V(r). \tag{15}$$

In that case a change δ in the barrier multiplies the mean lifetime by $e^{N\delta+o(N)}$, so even a modest change in ΔV has a large effect when N is large. Section 5.2 states which stopping times are computed.

The extinction time (2) counts every failed attempt from the saddle region. If successive attempts were independent, the number of attempts would be geometric with success probability p_N, and the mean extinction time would be the mean duration of an attempt divided by p_N. A success probability with $\log p_N = o(N)$ would then change the mean extinction time only by a subexponential factor, leaving the exponent in (15) unchanged. We do not prove that the jump process behaves in this way; Section 5.2 reports the success probabilities observed in simulation.

4.1 Barrier after eliminating the signal

Eliminating the signal removes the extra freedom in w. By (7), it imposes $w = \bar{w}(x) = \alpha_C x/\kappa$ and leaves a one-dimensional birth-death process for x, with birth rate $x\,\bar{b}(x)$ and death rate $x\,d(x)$. The eliminated-signal Hamiltonian is then

$$\mathbb{H}_\infty(x,p) = x\big[\bar{b}(x)(e^{p}-1) + d(x)(e^{-p}-1)\big],$$

i.e., the full Hamiltonian (12) at $w = \bar{w}(x)$ and $q = 0$, where $\mathbb{H}_f$ vanishes. Notice that the signal itself has not been switched off: its quasi-steady value still enters the cell division rate through $\bar{b}(x) = b(\bar{w}(x))$. We denote the barrier of this model by ΔV_∞. This is a definition, not an interchange of limits: the relation $w = \bar{w}(x)$ is the limit $r \to \infty$ taken after $N \to \infty$, and whether $\Delta V_\infty = \lim_{r\to\infty} \Delta V(r)$ is a separate question, left for [6].

The advantage of the eliminated-signal model is that its barrier can be calculated explicitly as follows. Along a free-time minimizing orbit the Hamiltonian is conserved and has value zero. Setting $z = e^{p} > 0$, from $\mathbb{H}_\infty(x,p) = 0$ and $x > 0$ we obtain

$$0 = \bar{b}(x)(z-1) + d(x)(z^{-1}-1), \qquad 0 = (z-1)\big(\bar{b}(x)z - d(x)\big).$$

There are therefore two zero-energy branches

$$p = 0, \qquad p = p^*(x) := \log\frac{d(x)}{\bar{b}(x)}. \tag{16}$$

Their directions follow from Hamilton's equation $\dot{x} = \partial_p \mathbb{H}_\infty$

$$\dot{x}\big|_{p=0} = x\big(\bar{b}(x) - d(x)\big) > 0, \qquad \dot{x}\big|_{p=p^*} = x\big(d(x) - \bar{b}(x)\big) < 0 \quad \text{for } x \in (x^*, x_{\mathrm{on}}).$$

The first branch is the deterministic motion back toward the cooperative state. The second branch moves toward the threshold and is the instanton of the eliminated-signal model.

Because $\mathbb{H}_\infty = 0$ on this branch, its action reduces to the line integral of the momentum

$$\begin{aligned}\Delta V_\infty &= \int p^*(x)\,dx = \int_{x_{\mathrm{on}}}^{x^*} \log\frac{d(x)}{\bar{b}(x)}\,dx \\ &= \int_{x^*}^{x_{\mathrm{on}}} \log\frac{b(\bar{w}(x))}{d(x)}\,dx.\end{aligned} \tag{17}$$

Notice that on (x^*, x_{on}) we have $\bar{b}(x) > d(x)$, hence the last integrand is positive, i.e., the instanton moves against the deterministic drift. At $c = 0.36$,

$$\Delta V_\infty = 0.06138.$$

4.2 What changes at finite r

For finite r, the zero-energy condition is

$$\mathbb{H}(x, w, p, q) = 0.$$

This is one equation for the two momenta p and q, so it no longer selects a unique momentum as a function of the state. Where the quasipotential is differentiable it still satisfies $\mathbb{H}(U, \nabla V_r^Q(U)) = 0$, and along a smooth minimizing path $(p, q) = \nabla V_r^Q$; but this Hamilton-Jacobi equation in two variables does not reduce to a choice between branches as in (16). The cell and signal coordinates also remain coupled through the coefficients of (12). These are the reasons why no one-dimensional quadrature analogous to (17) is available and why the finite-r problem must be minimized numerically in Section 5.1. The network is not reversible, which is why the minimizing path need not be the deterministic relaxation path run backwards (Section 4); it does not by itself prevent a potential.

Nevertheless, the eliminated-signal instanton gives a useful comparison path for large r. Since $\mathbb{H}_s$ contains only p and $\mathbb{H}_f$ only q, the Legendre transform separates

$$\mathbb{L}(U, v) = \mathbb{L}_s(U, v_1) + r\,\mathbb{L}_f\Big(U, \frac{v_2}{r}\Big).$$

Let $x(t)$ be the eliminated-signal instanton,

$$\dot{x} = x\,[d(x) - \bar{b}(x)], \tag{18}$$

and lift it to the plane by setting $U_{\mathrm{sl}}(t) = (x(t), \bar{w}(x(t)))$. Its cell contribution is exactly ΔV_∞; the remaining one is the cost of making the signal follow the slaved relation at a nonzero speed. Along this path the signal production and removal rates are equal, $a := \alpha_C x = \kappa \bar{w}(x)$, hence $\mathbb{H}_f = 2a(\cosh q - 1)$ and $\mathbb{L}_f(U, u) = u^2/(4a) + O(u^4)$. With $\dot{w} = (\alpha_C/\kappa)\dot{x}$ and, by (18), $\int \dot{x}^2/x\,dt = \int_{x^*}^{x_{\mathrm{on}}} [\bar{b}(x) - d(x)]\,dx$, we obtain

$$r\mathbb{L}_f\Big(U, \frac{\dot{w}}{r}\Big) = \frac{1}{r}\,\frac{\alpha_C \dot{x}^2}{4\kappa^2 x} + O(r^{-3}).$$

Hence the lifted path has action

$$\mathcal{S}(U_{\mathrm{sl}}) = \Delta V_\infty + \frac{C_{\mathrm{sl}}}{r} + O(r^{-3}), \qquad C_{\mathrm{sl}} = \frac{\alpha_C}{4\kappa^2}\int_{x^*}^{x_{\mathrm{on}}} \big[\bar{b}(x) - d(x)\big]\,dx. \tag{19}$$

The instanton reaches the fixed points only at infinite times. Truncating it near both endpoints and then letting the endpoint corrections shrink gives an admissible sequence for (13), so (19) implies

$$\Delta V(r) \le \Delta V_\infty + \frac{C_{\mathrm{sl}}}{r} + O(r^{-3}). \tag{20}$$

At $c = 0.36$, $C_{\mathrm{sl}} = 0.06579$.

The upper bound keeps the signal on the slaved relation (7). Consider a displacement of order $1/r$, i.e., $w = \bar{w}(x) + \xi/r$. The signal drift is $-\kappa\xi$, the cell Lagrangian changes by $(\xi/r)\,m$ with $m := \partial_w \mathbb{L}_s\big((x, \bar{w}(x)), \dot{x}\big)$, and expanding $\mathbb{L}_f$ about its drift as above we obtain, formally,

$$\mathcal{S}(U) = \int \mathbb{L}_s\big((x, \bar{w}(x)), \dot{x}\big)\,dt + \frac{1}{r}\int \Big(\xi\, m + \frac{\big(\bar{w}'(x)\dot{x} + \kappa\xi\big)^2}{4\alpha_C x}\Big)dt + O(r^{-2}). \tag{21}$$

Notice that $\dot{\xi}$ does not occur, hence the second integral is minimized pointwise in ξ, and to this order it is evaluated on the instanton (18), where $m = b'(\bar{w}(x))\bar{F}(x)/\bar{b}(x)$ by the envelope theorem. The choice $\xi = 0$ returns C_{sl}. A signal which follows its own drift, i.e., $\kappa\xi = -\bar{w}'(x)\dot{x}$, pays no signal action and gives the coefficient $A_{\mathrm{lag}} = (\alpha_C/\kappa^2)\int_{x^*}^{x_{\mathrm{on}}} m\,dx$. The minimizer gives

$$A = \frac{\alpha_C}{\kappa^2}\int_{x^*}^{x_{\mathrm{on}}}\left[m - \frac{x\,m^2}{\bar{F}(x)}\right]dx = \frac{\alpha_C}{\kappa^2}\int_{x^*}^{x_{\mathrm{on}}}\frac{b'(\bar{w}(x))\bar{F}(x)}{b(\bar{w}(x))}\left[1 - \frac{x b'(\bar{w}(x))}{b(\bar{w}(x))}\right]dx, \tag{22}$$

i.e., the lag of the signal raises the barrier by A_{lag}/r, and the fluctuations of the signal remove the second term. At $c = 0.36$ we have $C_{\mathrm{sl}} = 0.0658$, $A_{\mathrm{lag}} = 0.0472$ and $A = 0.0354$. If the signal count is scaled as νN, i.e., the signal jumps have size $1/(\nu N)$ and their rates are multiplied by ν, then the quadratic term in (21) is multiplied by ν and the same minimization gives

$$A(\nu) = A_{\mathrm{lag}} - \frac{A_{\mathrm{lag}} - A}{\nu}, \tag{23}$$

which interpolates between the present model, $\nu = 1$, and a signal without fluctuations, $\nu \to \infty$.

This calculation is formal. It assumes that the minimizing paths have the form $w = \bar{w}(x) + \xi/r$ uniformly in time, and it exchanges the minimization with the expansion; the details are given in [6]. Here we use (22) and (23) as predictions, which are tested in Section 6. The comparison path proves only the upper bound (20). Neither it nor the formal calculation proves rigorously that $\Delta V(r) \geq \Delta V_\infty$, nor that the leading correction is A/r.

5 Computing the barrier: full case

We develop the two routes in turn.

5.1 Minimum action

Discretization

We compute (13) by a minimum action method [10]. For a horizon T and a uniform mesh $t_k = k\Delta t$ with $\Delta t = T/M$, we represent a path by its nodes $\boldsymbol{U} = (U_0, \dots, U_M)$ with $U_0 = U_{\mathrm{on}}$ and $U_M = U_*$ held fixed, and approximate its action by

$$\mathcal{S}^{(r)}_{T,M}(\boldsymbol{U}) := \sum_{k=0}^{M-1} \Delta t\, \mathbb{L}^{(r)}\left(\frac{U_k + U_{k+1}}{2},\ \frac{U_{k+1} - U_k}{\Delta t}\right), \tag{24}$$

i.e., a midpoint rule in the state variable with forward-difference velocities. The interior nodes $U_1, \dots, U_{M-1}$ are the optimization variables. Since the endpoint $U_M = U_*$ is imposed, the method approximates the saddle barrier (13). It does not select an endpoint, and in particular it does not locate the cheapest point of the threshold line $\{x = x^*\}$.

Evaluating the Lagrangian

The implementation evaluates the Legendre transform pointwise. For a pair of opposite unit jumps with positive intensities λ_+ and λ_-, write

$$p^*(v; \lambda_+, \lambda_-) = \log\frac{v + \sqrt{v^2 + 4\lambda_+\lambda_-}}{2\lambda_+}, \qquad \ell(v; \lambda_+, \lambda_-) = vp^* - \sqrt{v^2 + 4\lambda_+\lambda_-} + \lambda_+ + \lambda_-. \tag{25}$$

The reaction directions in Fig. 2 are axis-aligned, hence the local rate is explicitly

$$\mathbb{L}(U,v) = \ell\big(v_x; xb(w), xd(x)\big) + \ell\big(v_w; r\alpha_C x, r\kappa w\big), \tag{26}$$

where the two intensities are passed as the second and third arguments of ℓ. This closed form is special to the present network.

Implementation

For transfer to networks with coupled jump directions, the production code computes the maximizing momentum by a generic damped Newton solve of

$$v_k = \partial_p \mathbb{H}^{(r)}(U_k, p_k), \qquad \mathbb{L}^{(r)}(U_k, v_k) = \langle p_k, v_k\rangle - \mathbb{H}^{(r)}(U_k, p_k),$$

which agrees with (25) and (26) to 4×10^{-15} on 2000 random state-velocity pairs. By the envelope theorem [25], $\partial_v \mathbb{L}^{(r)} = p^*$ and $\partial_U \mathbb{L}^{(r)} = -\partial_U \mathbb{H}^{(r)}$ at that momentum, so gradients need no further Legendre solves. The interior nodes are optimized by L-BFGS-B [8, 30] in the implementation of [28], with `ftol`$= 10^{-14}$, `gtol`$= 10^{-10}$ and trial states kept above 10^{-9}, starting from the straight segment from U_{on} to U_*.

Calibration and convergence

As an end-to-end calibration, we apply the same discretization and optimizer to the one-dimensional eliminated-signal problem. With the production settings $T = 40$, $M = 2000$, its action is 0.0613801, against $\Delta V_\infty = 0.0613795$ from (17), a relative error of 0.0010%.

That calibration exercises the one-dimensional eliminated-signal problem only. A second one exercises the two-coordinate minimizer itself, at finite r, against a closed form the solver never uses. Near an asymptotically stable equilibrium the rate functional of Theorem 5 gives a quadratic quasipotential, i.e., $V(U) = \frac{1}{2}\langle U - U_{\mathrm{on}}, \Pi(U - U_{\mathrm{on}})\rangle + O(|U - U_{\mathrm{on}}|^3)$, whose Hessian Π solves $\Pi J_r + J_r^\top \Pi + \Pi D \Pi = 0$, with D as in (11). Multiplying that identity by Π^{-1} on both sides returns the Lyapunov equation of Theorem 4, hence

$$\Pi = \Sigma^{-1}, \tag{27}$$

i.e., the quasipotential Hessian is the inverse stationary covariance of the Gaussian regime, available from a 2×2 Lyapunov solve. The identity (27) does not presume reversibility, which the network lacks (Kolmogorov's criterion fails around every unit cell of the lattice): it is an algebraic consequence of the Riccati and Lyapunov equations above. We therefore minimize the action from U_{on} to $U_{\mathrm{on}} + \varepsilon z$ over six unit directions z and two excursion sizes, extrapolate the quadratic forms $2\varepsilon^{-2}\mathcal{S}$ to $\varepsilon \to 0$, and obtain Π by least squares. Its three independent entries agree with (27) to a relative error of at most 2.5×10^{-4} over $r \in \{0.5, 1, 2, 4, 8\}$. Notice that this check shares nothing with the one-dimensional calibration and nothing with the exit-time experiment of Section 5.2: it tests the finite-r two-coordinate minimizer against Theorem 4. It is nevertheless local, hence it constrains the solver near U_{on} and says nothing about the global minimum in (13).

We separate horizon and mesh effects at the smallest and largest values of r. At fixed $\Delta t = 0.02$ the $r = 0.5$ barrier changes by 7.3×10^{-5} between $T = 20$ and $T = 30$ and by at most 10^{-6} up to $T = 60$, while $\Delta V(64)$ does not change; at fixed $T = 40$, refining Δt from 0.08 to 0.01 changes $\Delta V(0.5)$ by 3×10^{-6} and $\Delta V(64)$ by 10^{-5}, and the last halving by at most 10^{-6}. Both studies bracket the production setting $T = 40$, $M = 2000$.

The production sweep uses $T = 40$, $M = 2000$ and an outer iteration budget of 60000. The default budget of 600 iterations fails upward, returning 0.128412 at $r = 0.5$ instead of 0.122477. We compute $r \in \{0.5, 1, 2, 4, 8, 16, 32, 64\}$ for $c \in \{0.20, 0.30, 0.36, 0.45, 0.50\}$, in 234 seconds for all 40 minimizations; code, settings and barriers are in the accompanying repository [5].

The same discretization and optimizer are used for further sweeps at $c = 0.36$. For the slow-signal regime we compute nine rates between $r = 0.354$ and $r = 0.0078$ with the horizon $T = 20/r$, since the signal relaxes on the time $1/(r\kappa)$. Replacing it by $T = 30/r$ changes these barriers by at most 6×10^{-8}, refining Δt from 0.08 to 0.02 by at most 3×10^{-6}, and fresh starts from paths at $r = 0.125$ and $r = 0.0625$ agree to 1.2×10^{-6}. For the comparison with the diffusion approximation we replace $\mathbb{L}$ by its quadratic part $\frac{1}{2}\langle v - F, D^{-1}(v - F)\rangle$ of Remark 6; on the eliminated-signal problem this solver returns 0.0613218 against the closed form $\int_{x^*}^{x_{\mathrm{on}}} 2(\bar{b} - d)/(\bar{b} + d)\, dx = 0.0613211$. Moreover, for a signal count scaled as νN the signal term of (26) is multiplied by ν, and we compute $\nu \in \{1, 2, 4, 16, 64\}$ at $r \in \{8, 16, 32, 64\}$.

Three further checks address the global and the discrete character of the minimum at $c = 0.36$ and $r \in \{0.5, 1, 4, 16, 64\}$. First, besides the straight segment we start from the same segment bowed up and down in w and from the lifted eliminated-signal instanton of Section 4. All four starts converge to the same path, within Hausdorff distance 4×10^{-3}, and to the same action within 2.2×10^{-7}. Second, a free-time minimizer has $\mathbb{H} = 0$ along the path, and the computed paths give $\max_k |\mathbb{H}(U_k, p_k)| \le 9 \times 10^{-7}$. Third, the midpoint sum (24) is not itself an upper bound for the infimum, whereas the exact action of any admissible path is. The exact action of the piecewise-linear interpolant of each computed path, integrated by Gauss-Legendre quadrature on every interval, lies 7×10^{-8} to 4×10^{-7} below the midpoint value, hence the reported barriers are upper bounds for (13), and none of the starts reveals a lower minimum.

Extraction of the coefficient

For each displayed value of c we extract the coefficient $A_{\mathrm{num}}(c)$ from the large-r numerical barriers through

$$\Delta V_{\mathrm{num}}(r; c) = \Delta V_\infty(c) + \frac{A_{\mathrm{num}}(c)}{r} + O(r^{-2}), \tag{28}$$

with $\Delta V_\infty(c)$ fixed at its quadrature value, i.e., $A_{\mathrm{num}}(c)$ is the plateau of $r[\Delta V_{\mathrm{num}}(r; c) - \Delta V_\infty(c)]$ at large r. We read it at $r = 64$ and use the difference from $r = 32$ as a residual estimate. We compare it with the prediction (22), evaluated by quadrature at each cost; the formal calculation behind it is not used in the minimization.

5.2 Exit times

The minimum-action calculation and the expansion (3) share the same variational structure, hence agreement between them would not be an independent test. We therefore compute mean exit times directly from the generator (4), by solving its backward equation on a truncated lattice and by simulating the jump process. The primary stopping time is the extinction time (2). We also compute the first time the cell density reaches the threshold,

$$\tau_N^{\mathrm{thr}}(r) := \inf\{t \ge 0 : X_t^N \le x^*\}, \tag{29}$$

where X_t^N denotes the cell density. For a process started above the threshold, $\tau_N^{\mathrm{thr}}(r) \le \tau_N^{\mathrm{off}}(r)$, since $a_{\mathrm{off}} < x^*$ and the density moves by jumps of size $1/N$. Assuming metastable exit asymptotics [15], we write, for either stopping time,

$$\mathbb{E}[\tau_N(r)] = C_r(N) \exp\{N \Delta V_\tau(r)\}, \qquad N^{-1} \log C_r(N) \to 0, \tag{30}$$

where $\Delta V_\tau = \Delta V$ for the extinction time by (15) and $\Delta V_\tau = \Delta V_{\mathrm{thr}}$, defined below, for the threshold time. We do not prove (30) for this jump process, whose rates vanish on the axes; the computations below test it for $40 \le N \le 400$. For exit through an interior saddle the prefactor $C_r(N)$ is expected to tend to a constant, and this is indeed what we find. Notice that neither route constructs an instanton, minimizes an action, or uses a boundary-layer ansatz.

Backward equation

The mean $\tau(u) := \mathbb{E}_u[\tau_N]$ of either stopping time solves the backward equation $\mathfrak{L}_N \tau = -1$ outside the target, with $\tau = 0$ on the target. We solve it at $c = 0.36$ on the lattice truncated at U_{on} plus nine stationary standard deviations of Σ/N in each coordinate, where the jumps out of this box are suppressed. Since τ is of order $e^{N\Delta V}$, a direct solve loses accuracy as N grows, hence we use a renewal identity. Let u_0 be the lattice point nearest to U_{on}, which is also the starting state of the simulations below, with total jump rate q_0 and jump probabilities $P(u_0, j)$. For $u \neq u_0$ let $h(u)$ be the probability of reaching the target before u_0, and $\psi(u)$ the mean time to reach the target or u_0. By the strong Markov property at the returns to u_0,

$$\mathbb{E}_{u_0}[\tau_N] = \frac{q_0^{-1} + \sum_j P(u_0, j)\, \psi(j)}{\sum_j P(u_0, j)\, h(j)}. \tag{31}$$

The functions h and ψ solve two sparse linear systems on the lattice without u_0, which we solve directly. We use $N = 40, \ldots, 400$, with 25 to 31 sizes per rate including every simulated size, and three targets: R_{off}, the same rectangle with both sides halved, and the threshold (29). These are 399 solves with up to 1.4×10^6 unknowns, in 36 minutes of computing time on a standard laptop.

We checked the solver as follows. Enlarging the box from 9 to 12 standard deviations changes $\mathbb{E}[\tau_N]$ by at most 4×10^{-14} in relative terms, and a second elimination ordering by at most 3×10^{-14}, in the 15 cases $N \in \{40, 80, 120\}$, $r \in \{0.5, 1, 2, 4, 8\}$. The relative residual of the system for h is at most 9×10^{-15}. Where $N\Delta V \le 21$ the direct solve of $\mathfrak{L}_N \tau = -1$ agrees with (31) to 10^{-8} for $\tau < 10^7$ and drifts to 2.5×10^{-5} at the largest τ.

These values carry no sampling error. For each r we fit them, by least squares over $N \ge N_{\min}$, with

$$\log \mathbb{E}[\tau_N(r)] = a_r + N\Delta V_{\mathrm{BE}}(r) + \frac{c_r}{N} + \frac{e_r}{N^2}, \tag{32}$$

i.e., a constant prefactor e^{a_r} with its first two corrections, and we read the model error from the change of $\Delta V_{\mathrm{BE}}(r)$ with $N_{\min} \in \{120, 200\}$ and with $e_r = 0$.

Simulation protocol and estimator

With the help of the stochastic simulation algorithm [2, 16] we simulate the four channels of Fig. 2 exactly. Each replica starts from the integer state $n^C(0) = [N x_{\mathrm{on}}]$, $n^{\mathrm{w}}(0) = [N \bar{w}(x_{\mathrm{on}})]$, which depends on N but not on r. We use $r \in \{0.5, 1, 2, 4, 8\}$ at $c = 0.36$, with five to nine population sizes per rate in $40 \le N \le 140$, and each pair (r, N) has its own reproducibly seeded random stream. A first campaign, with 1500 to 2500 replicas per pair, stops at τ_N^{thr}. A second campaign follows 1500 replicas per pair until entry into $R_{\mathrm{off}} = [0, 0.3] \times [0, 0.7]$, which satisfies (8) at every cost used, and records τ_N^{thr} on the way. It uses a separately written simulator, checked against five stored points of the first campaign (combined $p = 0.20$).

Replicas are run to a deterministic censoring time $T^{\mathrm{c}}_{r,N}$. For replica k let t_k be the smaller of $\tau_{N,k}(r)$ and $T^{\mathrm{c}}_{r,N}$, and let $\delta_k = 1$ if $\tau_{N,k}(r) \le T^{\mathrm{c}}_{r,N}$ and $\delta_k = 0$ otherwise. Under the exponential

exit-time approximation, the right-censored maximum-likelihood estimator of the escape rate is given by

$$\widehat{\lambda}_{r,N} := \frac{\sum_k \delta_k}{\sum_k t_k}, \qquad \widehat{\mathbb{E}[\tau_N(r)]} := \widehat{\lambda}_{r,N}^{-1}.$$

The censoring time is set from the minimum-action barrier and no prefactor, i.e., $T^{\mathrm{c}}_{r,N} = \eta\, e^{N\Delta V_{\mathrm{num}}(r)}$, with $\eta = 30$ for the extinction time, since $\mathbb{E}[\tau_N^{\mathrm{off}}]$ is 3 to 4 times $\exp\{N\Delta V_{\mathrm{num}}(r)\}$ at these sizes, and $\eta = 1.5$ for the threshold time (0.3 at $r = 0.5$ with $N \geq 90$). Then 13 of the 51 000 extinction times are censored, and the threshold escape fractions range from 0.53 to 1.00. The action value sets the computational budget but is not inserted into the estimated mean. For each r, by weighted least squares over the simulated sizes, we fit

$$\log \widehat{\mathbb{E}[\tau_N(r)]} = a_r + N\Delta V_{\mathrm{SSA}}(r) + \frac{c_r}{N}, \tag{33}$$

with weights $n_{\mathrm{exit}}(r, N)$, since the variance of $\log \widehat{\mathbb{E}[\tau_N(r)]}$ is approximately n_{exit}^{-1}, and uncertainties are inflated by $\max(1, \chi^2_{\mathrm{red}})^{1/2}$. The term c_r/N is the first correction to a constant prefactor. Since the weights are known counting variances, the reference distribution of a standardized difference is normal when $\chi^2_{\mathrm{red}} \leq 1$. We nevertheless quote two-sided p-values from Student's t distribution with the degrees of freedom of the fit, which are the larger ones.

The estimator behaves as the exponential approximation requires. Over five independent seeds at $r = 0.5$, $N = 60$, the between-seed standard deviation of $\log \widehat{\mathbb{E}[\tau_N^{\mathrm{thr}}]}$ is 0.0236, against the nominal $1500^{-1/2} = 0.0258$, and censoring times imposed post hoc change the censored estimate by at most 2% at escape fractions near 0.65, with no systematic trend. At all 34 pairs of the second campaign the extinction times satisfy the exponential identities $\mathrm{sd}/\mathbb{E} = 1$ and $\mathrm{median}/\mathbb{E} = \log 2$ within sampling error: the ratios lie in 0.946–1.042 and 0.653–0.730, and the Kolmogorov-Smirnov test against the exponential law with the fitted mean, censored at $T^{\mathrm{c}}_{r,N}$, gives $p \geq 0.035$ at every pair, which is not more extreme than expected among 34 tests.

The extinction time

Figure 4 shows the prefactor $\mathbb{E}[\tau_N^{\mathrm{off}}]\, e^{-N\Delta V_{\mathrm{num}}(r)}$ of the backward equation against $1/N$. It decreases monotonically, from 3.3–4.2 at $N = 40$ to 2.5–3.3 at $N = 400$, and it is linear in $1/N$ at large N, i.e., $C_r(N)$ tends to a constant with a correction of relative order $1/N$. The fit (32) gives the growth rates $\Delta V_{\mathrm{BE}}(r)$ of Table 2. They differ from $\Delta V_{\mathrm{num}}(r)$ by at most 2.4×10^{-6} for $N_{\min} = 200$ and 6.4×10^{-6} for $N_{\min} = 120$, and by at most 8.5×10^{-6} with $e_r = 0$ and $N_{\min} = 200$. The model-free slopes of $\log \mathbb{E}[\tau_N^{\mathrm{off}}]$ between $N = 380$ and $N = 400$ lie 0.05% to 0.11% below $\Delta V_{\mathrm{num}}(r)$ and still increase with N. Halving both sides of the rectangle changes $\mathbb{E}[\tau_N^{\mathrm{off}}]$ by at most 2.2% at $N = 40$ and by less than 0.05% for $N \geq 100$. This is expected, since R_{off} lies deep in the off basin, where arrival costs no additional action.

The simulated means agree with these values. Over the 34 pairs (r, N) of the second campaign, the standardized differences between $\log \widehat{\mathbb{E}[\tau_N^{\mathrm{off}}]}$ and its exact value have mean 0.10 and standard deviation 0.97, the largest is 2.7 in absolute value, and their sum of squares is 31.3 ($p = 0.60$ for χ^2 with 34 degrees of freedom). Hence the solver and the simulator, which share no code, validate each other. Applied to the simulated means, the regression (33) gives the slopes $\Delta V^{\mathrm{off}}_{\mathrm{SSA}}(r)$ of Table 2, with reduced chi-square between 0.40 and 1.66. They lie 1.3, 0.2, 0.6, −0.5 and 0.6 fitted standard errors from $\Delta V_{\mathrm{num}}(r)$ at $r = 0.5, 1, 2, 4, 8$, and they stay within 1.4 standard errors of it at every rate when the smallest size, or the two smallest, are removed. Without the smallest size, the simulations still exclude ΔV_∞ at $r = 0.5$ and 1 ($p \leq 0.001$), but not at $r = 2$, 4 and 8 ($p \geq 0.20$); there the exclusion of ΔV_∞ rests on the backward equation.

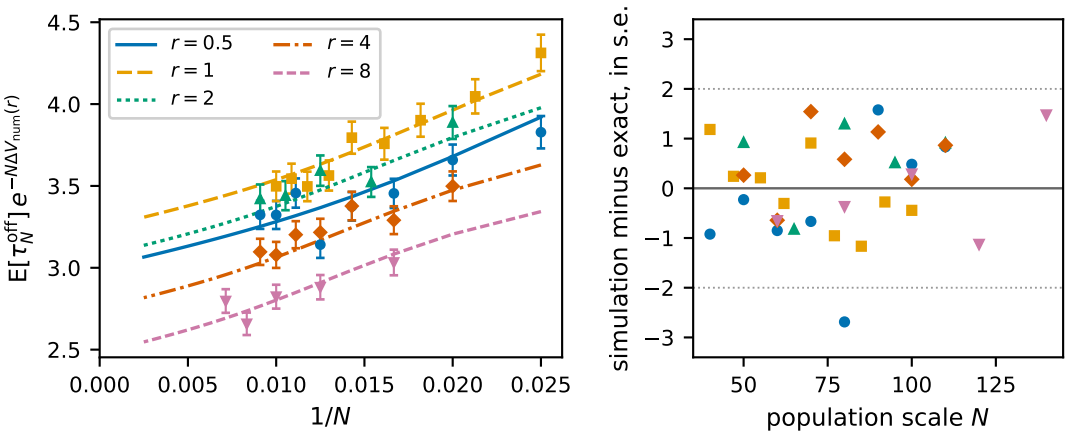


Figure 4: Mean extinction times (2) at $c = 0.36$. Left: the prefactor $\mathbb{E}[\tau_N^{\mathrm{off}}]\, e^{-N\Delta V_{\mathrm{num}}(r)}$ against $1/N$; curves are the backward equation for $40 \le N \le 400$, and points are the simulated means with nominal $n_{\mathrm{exit}}^{-1/2}$ error bars. Right: the difference between the simulated and the exact $\log \mathbb{E}[\tau_N^{\mathrm{off}}]$, in units of the nominal standard error, with the same markers. All quantities are dimensionless.

The regression model matters. With $c_r = 0$ every slope lies below $\Delta V_{\mathrm{num}}(r)$, by 0.0012 to 0.0033, i.e., by 2.3 to 7.4 fitted standard errors. This bias is not statistical. Fitted to the exact means at the simulated sizes, with the simulation weights, the model with $c_r = 0$ returns $\Delta V_{\mathrm{num}}(r)$ minus 0.0017 to 0.0027, and (33) returns it within 4×10^{-4}. The second campaign required about 1.3×10^5 core-seconds.

A cheaper threshold point and recovery

The stopping rule (29) allows every signal value, whereas (13) fixes both coordinates. The corresponding variational quantity is the interior threshold barrier $\Delta V_{\mathrm{thr}}(r) := \inf_{w>0} V_r^Q((x^*, w)) \le \Delta V(r)$. By Proposition 7 and the remark after it, the points (x^*, w) with $w < w^*$, which lie in the off basin, are no cheaper than the saddle; the points with $w > w^*$ lie in the cooperative basin, and whether one of them is cheaper has to be checked. We computed the profile $w \mapsto V_r^Q((x^*, w))$ with the method of Section 5.1, the endpoint now fixed at (x^*, w), at $r \in \{0.5, 1, 2, 4, 8\}$ for $w^* < w \le w^* + 5$, from several initial paths, with $T \in \{40, 80, 160\}$ and $\Delta t \in \{0.04, 0.02, 0.01\}$, the paths being constrained to $x \ge x^*$ before the endpoint (which does not change the infimum over the line, since $\mathbb{L} \ge 0$). At $r = 1, 2, 4, 8$ the computed profile increases in w and no computed point costs less than the saddle, so there we take the computed threshold barrier to be $\Delta V_{\mathrm{num}}(r)$; this is a computed agreement at these rates, not a statement for all $r \ge 1$. At $r = 0.5$ the profile has an interior minimum at $w = 1.808$, i.e., 0.69 above w^*, with action 0.11308 against 0.12248 at the saddle. The difference, 9.4×10^{-3}, compares with a variation below 10^{-6} under the refinements, and an independently written solver reproduces both values to 10^{-6}. The computed threshold value at $r = 0.5$ is therefore 7.7% below the computed saddle value, and we write $\Delta V_{\mathrm{thr}}(0.5) \le 0.1131$ for this computed upper estimate. At small r the cheapest threshold crossing keeps the signal high, and the deterministic flow then returns to U_{on}.

For the threshold time the stopping rule in counts is $X^N \le \lfloor Nx^* \rfloor / N$, i.e., the effective threshold lies $\{Nx^*\}/N$ below x^*, where $\{Nx^*\} = Nx^* - \lfloor Nx^* \rfloor$ oscillates with N. Accordingly,

we fit the exact threshold times of the backward equation by (32) with $e_r = 0$ and an additional term $g_r\{Nx^*\}$, over $N \geq 200$ (Table 3). These fits are rougher than those of the extinction time: the residuals are 1.3×10^{-3} to 2.7×10^{-3}, and the slope still moves by a few 10^{-4} with $N_{\min}$. Nevertheless, at $r = 0.5$ the threshold time grows at the rate 0.1124, which is 0.6% below the computed upper estimate 0.1131 of $\Delta V_{\mathrm{thr}}(0.5)$ and 8% below the saddle value 0.1225. At $r \geq 1$ the computed threshold barrier is attained at the saddle, where the first variation of the action along the line vanishes, hence the lattice coefficient there is a finite-N effect; indeed it decreases with $N_{\min}$, $g_1 = 0.163, 0.152, 0.143$ for $N_{\min} = 40, 120, 200$, and it is consistent with the first crossings at $r = 1$, which have $w > w^*$ in 99% of trajectories. The simulated threshold means agree with the exact ones: over 34 pairs each, the standardized differences have mean -0.19 and standard deviation 0.75 in the first campaign, and 0.25 and 0.91 in the second. Simulations with the threshold shifted by whole lattice steps are kept in the repository [5].

Threshold crossings are often followed by recovery. Of all crossings, the fraction that returns to the box $[x_{\mathrm{on}} \pm 0.15] \times [\bar{w}(x_{\mathrm{on}}) \pm 0.30]$ before reaching R_{off} is 0.86–0.97 at $r = 0.5$ and 0.37–0.46 at $r = 8$, generally growing with N. At $r = 0.5$ the signal at the first crossing exceeds w^* in more than 99% of trajectories, with median excess between 0.68 and 0.83, close to the minimizing endpoint above; at $r = 8$ it does so in 81–83%. Such crossings lie in the cooperative basin (Section 3.1.2). At $r = 0.5$ the probability that a crossing with $w > w^*$ reaches R_{off} before returning falls from 0.14 at $N = 40$ to 0.03 at $N = 110$. Crossings with $w \leq w^*$, which lie in the off basin, still return in 15–40% of cases at every rate, with no clear trend in N.

Finally, we test whether the success probability near the saddle stays of order one, as the attempt argument of Section 4 requires. We start the process at the lattice point nearest to $U_* + sN^{-1/2}v_u$, where v_u is the unit unstable eigenvector at U_* and $s \in \{-1, 0, 1\}$, and stop at the first entry into R_{off} or into the box around U_{on}, with 20 000 trajectories for each $r \in \{0.5, 1, 2, 4, 8\}$ and $N \in \{50, 100, 200, \ldots, 3200\}$. Started at the saddle, the probability of entering R_{off} first decreases from 0.63–0.64 at $N = 50$ to 0.51–0.52 at $N = 3200$, with 95% intervals of half-width at most 0.007; at $N = 3200$ it is 0.74–0.79 for $s = -1$ and 0.23–0.29 for $s = 1$. Over this range it is therefore of order one, although the finite range of N cannot establish the limit.

6 Finite r raises the extinction barrier

Figure 5 shows the minimum-action results. The left panel exposes the mechanism: finite-r minimum-action paths bow above the quasi-steady signal relation and approach it as r increases. The middle panel shows the corresponding barriers. At $r = 1$ the barrier is 0.0953, i.e., it exceeds the eliminated-signal value $\Delta V_\infty = 0.0614$ by 55%, and between $r = 0.5$ and $r = 64$ it changes by a factor of 1.98, although the deterministic equilibria and their stability types are independent of r. The right panel tests the exponent in (3) without using the analytical coefficient: for each of five costs, $r[\Delta V_{\mathrm{num}}(r) - \Delta V_\infty]$ approaches a plateau. Over $r \geq 4$ the plateau spread is 0.57% at $c = 0.36$ and at most 3.5% across the five costs. The coefficients extracted through (28) agree with the formal prediction (22) to relative errors 0.56%, 0.36%, 0.24%, 0.08% and 0.01% as c ranges from 0.20 to 0.50.

The calculation of Section 4.2 separates the lag of the signal from its noise. At $c = 0.36$, a signal which follows its drift without fluctuations raises the barrier by A_{lag}/r with $A_{\mathrm{lag}} = 0.0472$, and the fluctuations of a signal as abundant as the cells remove 25% of this increase, leaving $A = 0.0354$. We test this by scaling the signal count as νN. At $r = 64$ the plateau $r[\Delta V_\nu(r) - \Delta V_\infty]$ equals 0.03553, 0.04140, 0.04433, 0.04653 and 0.04709 for $\nu = 1, 2, 4, 16, 64$, against $A(\nu) = 0.03544$, 0.04132, 0.04426, 0.04646 and 0.04701 from (23), i.e., a difference of

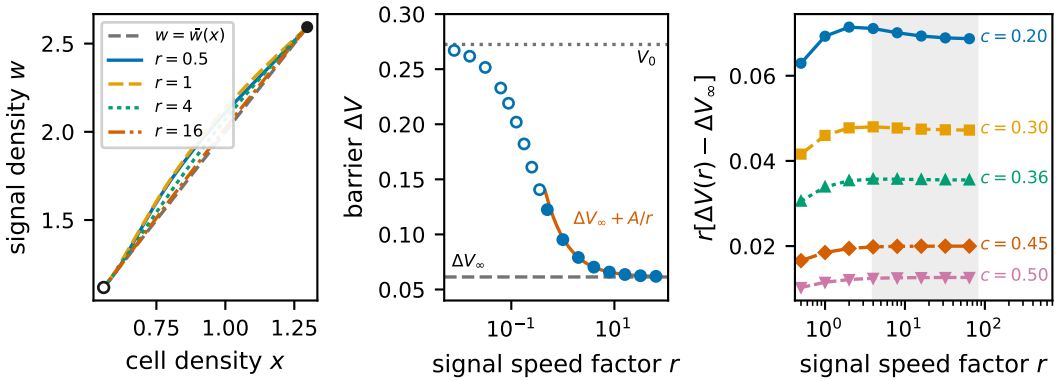


Figure 5: Minimum-action crossover. Left: minimum-action paths from U_{on} (filled point) to the saddle U_* (open point) at $c = 0.36$, with the endpoint imposed as in (13); the dashed line is the quasi-steady relation. Middle: finite-r barriers, with open points for the slow-signal sweep $r < 0.5$, the eliminated-signal value ΔV_∞, the level V_0 obtained by eliminating the cells, and $\Delta V_\infty + A/r$ using (22). Right: the rescaled excess barrier at five costs. A plateau identifies the $1/r$ scaling and its level gives A_{num}; the shaded region marks $r \geq 4$. All quantities are dimensionless.

0.16% to 0.24%, which is the size of the $O(1/r)$ remainder at $\nu = 1$. Hence the barrier is raised by the lag, and the signal noise offsets a part of it which vanishes as $\nu \to \infty$.

The expansion (3) concerns $r \to \infty$. Compared with the computed barriers, $\Delta V_\infty + A/r$ is 0.06% too large at $r = 2$, 1.6% at $r = 1$ and 8.0% at $r = 0.5$, and it diverges as $r \to 0$. For small r the roles of the two variables are exchanged. A formal elimination of the cells, $x = x(w) := (b(w) - d_0 - c)/\theta$, leaves a birth-death process for the signal, whose barrier is

$$V_0 = \int_{w^*}^{\bar{w}(x_{\mathrm{on}})} \log \frac{\alpha_C\, x(w)}{\kappa w}\, dw = 0.2724.$$

The open points in the middle panel of Fig. 5 are the slow-signal sweep of Section 5.1. The barrier increases monotonically as r decreases, from 0.1225 at $r = 0.5$ to 0.2670 at $r = 0.0078$, and $V_0 - \Delta V(r)$ is close to linear in r: the ratio $[V_0 - \Delta V(r)]/r$ equals 0.63, 0.67, 0.68 and 0.69 at $r = 0.0625$, 0.031, 0.016 and 0.0078. These computations support V_0 as the limit of $\Delta V(r)$ as $r \to 0$, with a correction of order r; we do not prove this. Over all the tested range of r the barrier therefore varies between 0.061 and 0.272, i.e., by a factor 4.4, with the equilibria fixed.

The barrier of the diffusion approximation, i.e., of the quadratic Lagrangian of Remark 6, lies 0.07% below $\Delta V_{\mathrm{num}}(r)$ at $r = 0.5$ and 0.10% to 0.11% below it at the seven rates from $r = 1$ to $r = 64$. Hence at these parameters the barrier, including its dependence on r, is a small-momentum quantity: the jump structure enters through the gross rates in D, and the higher cumulants contribute 0.1%. This concerns the exponent only, and it does not show that the diffusion approximation gives the right prefactor.

Table 2 compares the growth rates of the mean extinction time with the minimum-action saddle barrier. They test the prediction (15) directly, including the time spent in failed attempts. The rates $\Delta V_{\mathrm{BE}}(r)$ of the backward equation agree with $\Delta V_{\mathrm{num}}(r)$ to within 2.4×10^{-6} at all five rates, although the two computations share nothing: the first solves linear systems for the generator on the lattice, and the second minimizes the path action. The distance from $\Delta V_{\mathrm{num}}(r)$ to ΔV_∞ is 0.0611 at $r = 0.5$ and 0.0045 at $r = 8$, i.e., at least three orders of magnitude larger.

Hence the eliminated-signal value is excluded at every rate, including $r = 4$ and $r = 8$, where the two candidates differ by only 13% and 7%. The simulated slopes are consistent with $\Delta V_{\mathrm{num}}(r)$ at every rate. By themselves they exclude ΔV_∞ with $p \approx 10^{-6}$ or smaller at $r = 0.5, 1$, with $p = 0.016$ and $p = 0.022$ at $r = 2, 4$, and not at $r = 8$, where $p = 0.10$. Notice that all of this concerns $40 \le N \le 400$: the backward equation is a finite-N computation, and the limit in (15) remains an extrapolation. The minimization alone bounds each infimum from above, so it cannot establish that the barrier exceeds ΔV_∞; what excludes the eliminated-signal value is the backward equation, together with the simulation, without using the action solver. The minimum-action calculation supplies in return the mechanism and the $1/r$ law, neither of which the exclusion alone would reveal.

Table 2: Extinction barriers at $c = 0.36$. ΔV_{num}: minimum-action saddle barrier; ΔV_{BE} and $C_r(\infty)$: growth rate and limiting prefactor e^{a_r} of the mean extinction time of the backward equation, from (32) with $N_{\min} = 200$; $\Delta V_{\mathrm{SSA}}^{\mathrm{off}}$: slope of the simulated extinction times under (33), with its reduced chi-square and degrees of freedom; $c_r = 0$: the same fit without the term c_r/N. The final column gives the two-sided Student p-value of $\Delta V_{\mathrm{SSA}}^{\mathrm{off}} - \Delta V_\infty$, with $\Delta V_\infty = 0.061379$.

r	sizes	$\Delta V_{\mathrm{num}}(r)$	$\Delta V_{\mathrm{BE}}(r)$	$C_r(\infty)$	$\Delta V_{\mathrm{SSA}}^{\mathrm{off}}(r)$, χ^2_{red} (dof)	$c_r = 0$	p
0.5	8	0.122477	0.122477	3.00	0.1249 ± 0.0018, 1.66 (5)	0.1206 ± 0.0007	4×10^{-7}
1	9	0.095259	0.095258	3.25	0.0956 ± 0.0017, 0.40 (6)	0.0919 ± 0.0004	1×10^{-6}
2	5	0.079057	0.079056	3.08	0.0806 ± 0.0025, 1.06 (2)	0.0772 ± 0.0006	0.016
4	7	0.070313	0.070311	2.76	0.0693 ± 0.0022, 0.61 (4)	0.0683 ± 0.0005	0.022
8	5	0.065843	0.065840	2.49	0.0670 ± 0.0019, 1.35 (2)	0.0646 ± 0.0005	0.10

At $r = 0.5$ the extinction time does not follow the cheaper threshold point: its growth rate 0.12248 exceeds the computed threshold barrier $\Delta V_{\mathrm{thr}}(0.5) \le 0.1131$ by 0.0094, whereas the threshold time grows at the rate 0.1124 (Table 3). At $r = 1, 2, 4, 8$, where the computed threshold barrier equals the saddle value, the growth rates of the threshold time agree with it within 4×10^{-4}. These are finite-N findings, and they support the conditional prediction without proving it.

Table 3: Threshold barriers at $c = 0.36$. ΔV_{thr}: computed minimum-action threshold barrier of Section 5.2, equal to the saddle value for $r \ge 1$ and an upper estimate at $r = 0.5$; $\Delta V_{\mathrm{BE}}^{\mathrm{thr}}$ and g_r: growth rate of the mean threshold time of the backward equation and coefficient of the lattice term $g_r\{Nx^*\}$, fitted over $N \ge 200$; the last column is the root mean square residual of that fit.

r	$\Delta V_{\mathrm{thr}}(r)$	$\Delta V_{\mathrm{BE}}^{\mathrm{thr}}(r)$	g_r	residual
0.5	0.1131	0.1124	0.28	0.0014
1	0.0953	0.0957	0.14	0.0027
2	0.0791	0.0792	0.08	0.0022
4	0.0703	0.0703	0.06	0.0016
8	0.0658	0.0659	0.05	0.0013

The 55% excess and the coefficient A are parameter-specific. For this network, the $O(1/r)$ correction is observed at all five costs, and it is amplified exponentially in the population scale.

7 Conclusion

Within its stated numerical accuracy, the calculation establishes the central point: a parameter that leaves every deterministic equilibrium and stability type unchanged can alter the extinction barrier, and hence, subject to the exit estimates of Section 4, the mean extinction time by a factor exponential in the population size. Eliminating the fast signal before constructing the action removes the joint cell-signal history that produces this difference. Rare events are therefore not determined by the deterministic landscape alone, and a reduction that is harmless for the equilibria of a network need not be harmless for its exit times.

The two numerical routes are consistent, but they do not remove every source of uncertainty. The minimization of (24) is nonconvex, hence the value returned by the optimizer bounds the infimum in (13) from above, as the exact action of the interpolated paths confirms at the five rates checked, but it need not attain it. The calibrations of Section 5.1 are local, in the one-dimensional case and, through (27), at U_{on} at finite r, so neither verifies that the finite-r optimizer has found the global minimum. Proposition 7 identifies the saddle as the least-cost entrance to the basin of extinction, and the agreement with the threshold experiment at $r = 1, 2, 4, 8$ gives no indication of a cheaper route there; neither shows that the computed path is a global minimizer. At $r = 0.5$ the threshold computation shows that a different stopping set can have a smaller computed barrier. Three limitations therefore remain separate: global minimization of the action, the relation between the saddle barrier and stopping sets other than R_{off}, and the metastable exit estimates for the jump process, including the effect of the coordinate axes. For the extinction time itself the last is supported numerically by the backward equation and the simulations of Section 5.2, for $40 \le N \le 400$ only. Likewise the finite-T actions approach (13) from above, and Section 5.1 shows no change at the reported precision between $T = 40$ and $T = 60$, which controls the horizon truncation numerically but is not an analytical error bound.

Over $40 \le N \le 400$ the prefactor $C_r(N)$ in (30) tends to a constant between 2.5 and 3.3, with a correction of relative order $1/N$. A formula for this constant would require an Eyring-Kramers analysis for this non-reversible jump process, which, to our knowledge, is not available in the literature considered here. Likewise, the expansion (3) with the coefficient (22) and the limit V_0 at small r are supported numerically and not rigorously proved. Finally, the generic momentum inversion and envelope gradient extend directly to more channels and coordinates, but in higher dimensions additional saddles introduce distinct local minima, and such networks require multiple initial paths and, where necessary, a geometric minimum-action or string method [18].

A Proof of Proposition 7

Throughout, $c \in [0.2, 0.5]$ and $r > 0$ are fixed, Φ_t denotes the flow of (6), and $V = V_r^Q$. Recall from Section 3.1.2 that $b_0 < d_0 + c$, that $\bar{F}$ has exactly two positive zeros, both simple, and that the equilibria in $[0, \infty)^2$ are U_{off}, the hyperbolic saddle U_* and U_{on}. At U_* the Jacobian is

$$J_r = \begin{pmatrix} -\theta x^* & x^* b'(w^*) \\ r\alpha_C & -r\kappa \end{pmatrix},$$

with eigenvalues $\mu_s < 0 < \mu_u$. An eigenvector for μ is $(x^* b'(w^*), \mu + \theta x^*)$, which has two positive components for μ_u.

Lemma 8. *(i) Every solution started in Q converges to U_{off}, U_* or U_{on}. As a consequence, $\partial\mathcal{B}_{\mathrm{on}}(r) \cap Q \subset W^s(U_*)$.*

(ii) The branch of the unstable manifold of U_* that leaves towards $x < x^*$ lies in Q and converges to U_{off}.

Proof. (i) Let $u \in Q$. Since $\dot{x} \le x(b_0 + b_1 - \theta x)$ and $\dot{w} \le r(\alpha_C x - \kappa w)$, the solution is bounded, so its ω-limit set $\omega(u)$ is nonempty, compact, connected and invariant. If $\omega(u)$ contains a point $(0, w)$, it contains the solution through it, which converges to U_{off} along the axis; if it contains $(x, 0)$ with $x > 0$, the backward solution through that point leaves $[0, \infty)^2$, since $\dot{w} = r\alpha_C x > 0$ there, which is impossible. Since the stable equilibria attract a neighbourhood, $\omega(u)$ is then $\{U_{\text{off}}\}$, $\{U_{\text{on}}\}$, or a subset of Q that contains no equilibrium other than U_*. In the last case, with the Dulac function $g = 1/x$,

$$\operatorname{div}(gF) = \partial_x\big[b(w) - d(x)\big] + \partial_w\Big[\frac{r(\alpha_C x - \kappa w)}{x}\Big] = -\theta - \frac{r\kappa}{x} < 0$$

on Q, which is convex, hence Q contains neither a periodic orbit nor a homoclinic loop at U_*, and by the Poincaré-Bendixson theorem $\omega(u) = \{U_*\}$. Finally, the basins of the two stable equilibria are open and disjoint, so a point of $\partial\mathcal{B}_{\text{on}}(r) \cap Q$ lies in neither of them and converges to U_*.

(ii) For $0 \le x \le x^*$ we have $b_0 < d(0) \le d(x) \le d(x^*) = b(w^*) < b_0 + b_1$, hence $\zeta := b^{-1} \circ d$ is well defined, positive and C^1 on $[0, x^*]$, with $\zeta(x^*) = w^*$ and $\zeta' = \theta / b'(\zeta)$; the curve $w = \zeta(x)$ is the nullcline $\dot{x} = 0$. Since $\bar{F} < 0$ on $(0, x^*)$ and $b_0 < d_0 + c$, also $\bar{w}(x) < \zeta(x)$ on $[0, x^*)$. Let $\mathcal{O} := \{(x, w) : 0 < x < x^*,\ 0 < w < \zeta(x)\}$. The edge $x = 0$ is invariant; on the remaining open boundary arcs the field points inward: on $w = \zeta(x)$ we have $\dot{x} = 0$ and $\dot{w} = r\kappa[\bar{w}(x) - \zeta(x)] < 0$, on $x = x^*$ with $w < w^*$ we have $\dot{x} = x^*[b(w) - b(w^*)] < 0$, and on $w = 0$ we have $\dot{w} > 0$. Hence $\mathcal{O}$ is forward invariant. Inside $\mathcal{O}$, $\dot{x} < 0$. By (i), every trajectory in $\mathcal{O}$ converges to an equilibrium; since $x(t) \le x(0) < x^*$, that equilibrium must be U_{off}. Near U_* the unstable branch towards $x < x^*$ has slope $s_u = (\mu_u + \theta x^*)/(x^* b'(w^*))$, and $s_u - \zeta'(x^*) = \mu_u/(x^* b'(w^*)) > 0$. Hence it lies below the nullcline for $x < x^*$ close to x^*, i.e., in $\mathcal{O}$, and it converges to U_{off}. □

Lemma 9. *(i) $\mathbb{L}(U, v) \ge 0$ and $\mathbb{L}(U, F(U)) = 0$ for $U \in Q$ and $v \in \mathbb{R}^2$. As a consequence, $V(u, \Phi_t(u)) = 0$ for $u \in Q$ and $t > 0$.*

(ii) $V(u, z) \le V(u, y) + V(y, z)$ for $u, y, z \in Q$.

(iii) For every compact convex $K \subset Q$ there is $C_K < \infty$ such that $V(u, z) \le C_K |u - z|$ for $u, z \in K$.

Proof. (i) Taking $p = 0$ in the supremum gives $\mathbb{L} \ge 0$, since $\mathbb{H}(U, 0) = 0$. By convexity and $\nabla_p \mathbb{H}(U, 0) = F(U)$ (Remark 6), $\mathbb{H}(U, p) \ge \langle F(U), p\rangle$, hence $\mathbb{L}(U, F(U)) = 0$, and the deterministic solution, which stays in Q, has zero action. (ii) Concatenate two paths whose actions are within ε of $V(u, y)$ and $V(y, z)$, and let $\varepsilon \to 0$. (iii) On K the four intensities are continuous and bounded away from zero, hence by (25) and (26) the local rate $\mathbb{L}$ is bounded by some C_K on $K \times \{|v| \le 1\}$. The segment from u to z traversed at unit speed stays in K and has action at most $C_K |u - z|$. □

Proof of Proposition 7. Let $K \subset Q$ be a closed ball of radius $\rho > 0$ centred at U_*, and let C_K be as in Lemma 9(iii).

Points of the separatrix are no cheaper than the saddle. Let $y \in W^s(U_*) \cap Q$ and $0 < \varepsilon < \rho$, and choose $t > 0$ with $|\Phi_t(y) - U_*| \le \varepsilon$. By Lemma 9,

$$\Delta V(r) \le V(y) + V\big(y, \Phi_t(y)\big) + V\big(\Phi_t(y), U_*\big) \le V(y) + C_K \varepsilon,$$

and $\varepsilon \to 0$ gives $V(y) \geq \Delta V(r)$.

Lower bound. Let $z \in R \cap Q$ and let $U \in \mathrm{AC}([0,T];Q)$ run from U_{on} to z. Since U_{on} lies in the open set $\mathcal{B}_{\mathrm{on}}(r)$ and z does not, $t_1 := \inf\{t : U(t) \notin \mathcal{B}_{\mathrm{on}}(r)\}$ satisfies $0 < t_1 \leq T$, and $y := U(t_1) \in \partial\mathcal{B}_{\mathrm{on}}(r) \cap Q$, hence $y \in W^s(U_*)$ by Lemma 8(i). Since $\mathbb{L} \geq 0$,

$$\mathcal{S}_T(U) \geq \mathcal{S}_{t_1}\big(U|_{[0,t_1]}\big) \geq V(y) \geq \Delta V(r),$$

and taking the infimum over T, U and z gives $\inf_{z \in R\cap Q} V(z) \geq \Delta V(r)$.

Upper bound. Let Γ be the unstable branch of Lemma 8(ii). Given $0 < \varepsilon < \rho$, choose $u_\varepsilon \in \Gamma$ with $|u_\varepsilon - U_*| \leq \varepsilon$. Since R is a relative neighbourhood of U_{off}, there is $s > 0$ with $z_\varepsilon := \Phi_s(u_\varepsilon) \in R \cap Q$, and by Lemma 9

$$V(z_\varepsilon) \leq V(U_*) + V(U_*, u_\varepsilon) + V(u_\varepsilon, z_\varepsilon) \leq \Delta V(r) + C_K\varepsilon.$$

Hence $\inf_{z\in R\cap Q} V(z) \leq \Delta V(r) + C_K\varepsilon$ for every small $\varepsilon > 0$, and $\varepsilon \to 0$ gives the upper bound. □

References

[1] A. Agazzi, L. Andreis, R. I. A. Patterson, and D. R. M. Renger. Large deviations for Markov jump processes with uniformly diminishing rates. *Stoch. Process. Appl.*, 152:533–559, 2022.

[2] D. F. Anderson and T. G. Kurtz. *Stochastic Analysis of Biochemical Systems*. Mathematical Biosciences Institute Lecture Series. Springer, 2015.

[3] M. Assaf and B. Meerson. WKB theory of large deviations in stochastic populations. *J. Phys. A: Math. Theor.*, 50(26):263001, 2017.

[4] M. Assaf, E. Roberts, and Z. Luthey-Schulten. Determining the stability of genetic switches: explicitly accounting for mRNA noise. *Phys. Rev. Lett.*, 106(24):248102, 2011.

[5] M. Ayala and J. Zimmer. Computing extinction barriers in a quorum-sensing reaction network: code and data. Zenodo, version v1.1-mmsc2026, https://doi.org/10.5281/zenodo.23348482, 2026. Minimum-action and exit-time computations, solver settings, and the stored barriers reproducing the figures of this paper. Source repository https://github.com/marioayala15/qs-extinction-barriers.

[6] M. Ayala and J. Zimmer. Cooperation through a shared environment: establishment, spatial ecology, and rare collapse, 2026. Manuscript in preparation.

[7] F. Bouchet, T. Grafke, T. Tangarife, and E. Vanden-Eijnden. Large deviations in fast–slow systems. *J. Stat. Phys.*, 162(4):793–812, 2016.

[8] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu. A limited memory algorithm for bound constrained optimization. *SIAM J. Sci. Comput.*, 16(5):1190–1208, 1995.

[9] S. P. Diggle, A. S. Griffin, G. S. Campbell, and S. A. West. Cooperation and conflict in quorum-sensing bacterial populations. *Nature*, 450(7168):411–414, 2007.

[10] W. E, W. Ren, and E. Vanden-Eijnden. Minimum action method for the study of rare events. *Commun. Pure Appl. Math.*, 57(5):637–656, 2004.

[11] V. Elgart and A. Kamenev. Rare event statistics in reaction–diffusion systems. *Phys. Rev. E*, 70(4):041106, 2004.

[12] S. N. Ethier and T. G. Kurtz. *Markov Processes: Characterization and Convergence*, volume 282 of *Wiley Series in Probability and Statistics*. John Wiley & Sons, 2009.

[13] J. Feng and T. G. Kurtz. *Large Deviations for Stochastic Processes*, volume 131 of *Mathematical Surveys and Monographs*. American Mathematical Society, 2006.

[14] N. Fournier and S. Méléard. A microscopic probabilistic description of a locally regulated population and macroscopic approximations. *Ann. Appl. Probab.*, 14(4):1880–1919, 2004.

[15] M. I. Freidlin and A. D. Wentzell. *Random Perturbations of Dynamical Systems*, volume 260 of *Grundlehren der mathematischen Wissenschaften*. Springer, third edition, 2012.

[16] D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, 81(25):2340–2361, 1977.

[17] B. A. Hense, C. Kuttler, J. Müller, M. Rothballer, A. Hartmann, and J.-U. Kreft. Does efficiency sensing unify diffusion and quorum sensing? *Nat. Rev. Microbiol.*, 5(3):230–239, 2007.

[18] M. Heymann and E. Vanden-Eijnden. The geometric minimum action method: a least action principle on the space of curves. *Commun. Pure Appl. Math.*, 61(8):1052–1117, 2008.

[19] M. Khasin and M. I. Dykman. Extinction rate fragility in population dynamics. *Phys. Rev. Lett.*, 103(6):068101, 2009.

[20] T. G. Kurtz. Solutions of ordinary differential equations as limits of pure jump Markov processes. *J. Appl. Probab.*, 7(1):49–58, 1970.

[21] C. Léonard. Large deviations for long range interacting particle systems with jumps. *Ann. Inst. Henri Poincaré Probab. Stat.*, 31(2):289–323, 1995.

[22] M. B. Miller and B. L. Bassler. Quorum sensing in bacteria. *Annu. Rev. Microbiol.*, 55:165–199, 2001.

[23] J. M. Newby, P. C. Bressloff, and J. P. Keener. Breakdown of fast–slow analysis in an excitable system with channel noise. *Phys. Rev. Lett.*, 111(12):128101, 2013.

[24] R. I. A. Patterson and D. R. M. Renger. Large deviations of jump process fluxes. *Math. Phys. Anal. Geom.*, 22(3):21, 2019.

[25] R. T. Rockafellar. *Convex Analysis*. Number 28 in Princeton Mathematical Series. Princeton University Press, 1970.

[26] A. Shwartz and A. Weiss. *Large Deviations for Performance Analysis: Queues, Communication and Computing*. Chapman & Hall, 1995.

[27] P. Thomas, R. Grima, and A. V. Straube. Rigorous elimination of fast stochastic variables from the linear noise approximation using projection operators. *Phys. Rev. E*, 86(4):041110, 2012.

[28] P. Virtanen, R. Gommers, T. E. Oliphant, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods*, 17(3):261–272, 2020.

[29] C. M. Waters and B. L. Bassler. Quorum sensing: cell-to-cell communication in bacteria. *Annu. Rev. Cell Dev. Biol.*, 21:319–346, 2005.

[30] C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal. Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Trans. Math. Softw.*, 23(4):550–560, 1997.